



**Measuring Job-Finding Rates and Matching Efficiency
with Heterogeneous Jobseekers***

Robert E. Hall and Sam Schulhofer-Wohl

Economics Working Paper 15103

HOOVER INSTITUTION
434 GALVEZ MALL
STANFORD UNIVERSITY
STANFORD, CA 94305-6010

January 30, 2015

Matching efficiency is the productivity of the process for matching jobseekers to available jobs. Job-finding is the output; vacant jobs and active jobseekers are the inputs. Measurement of matching efficiency follows the same principles as measuring a Hicks-neutral index of productivity of production. We develop a framework for measuring matching productivity when the population of jobseekers is heterogeneous. The efficiency index for each type of jobseeker is the monthly job-finding rate for the type adjusted for the overall tightness of the labor market. We find that overall matching efficiency declined over the period, at just below its earlier downward trend. We develop a new approach to measuring matching rates that avoids counting short- duration jobs as successes. And we show that the outward shift in the Beveridge curve in the post-crisis period is the result of pre-crisis trends, not a downward shift in matching efficiency attributable to the crisis.

*The Hoover Institution supported Hall's research. The research is also part of the National Bureau of Economic Research's Economic Fluctuations and Growth Program. We thank Suyoun Han for excellent research assistance and Christopher Nekarda for sharing his Stata code for longitudinally matching observations in the Current Population Survey. The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

The Hoover Institution Economics Working Paper Series allows authors to distribute research for discussion and comment among other researchers. Working papers reflect the views of the authors and not the views of the Hoover Institution.

Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Jobseekers*

Robert E. Hall

Hoover Institution and Department of Economics,
Stanford University

National Bureau of Economic Research

rehall@stanford.edu; stanford.edu/~rehall

Sam Schulhofer-Wohl

Federal Reserve Bank of Minneapolis

wohls@minneapolisfed.org

January 30, 2015

Abstract

Matching efficiency is the productivity of the process for matching jobseekers to available jobs. Job-finding is the output; vacant jobs and active jobseekers are the inputs. Measurement of matching efficiency follows the same principles as measuring a Hicks-neutral index of productivity of production. We develop a framework for measuring matching productivity when the population of jobseekers is heterogeneous. The efficiency index for each type of jobseeker is the monthly job-finding rate for the type adjusted for the overall tightness of the labor market. We find that overall matching efficiency declined over the period, at just below its earlier downward trend. We develop a new approach to measuring matching rates that avoids counting short-duration jobs as successes. And we show that the outward shift in the Beveridge curve in the post-crisis period is the result of pre-crisis trends, not a downward shift in matching efficiency attributable to the crisis.

JEL E24, J63

*The Hoover Institution supported Hall's research. The research is also part of the National Bureau of Economic Research's Economic Fluctuations and Growth Program. We thank Suyoun Han for excellent research assistance and Christopher Nekarda for sharing his Stata code for longitudinally matching observations in the Current Population Survey. The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Matching efficiency is a key concept in understanding turnover in the labor market. In particular, turnover models imply that a decline in matching efficiency causes a rise in unemployment. Persistent high unemployment has generated concern that the U.S. economy's normal unemployment rate rose from the turmoil of the collapse of the housing market and the subsequent financial crisis. Similar concerns have developed in previous recessions.

The idea has proven useful that matching is a productive process that combines the efforts of jobseekers and of recruiting employers. The matching function—a central feature of the Diamond-Mortensen-Pissarides model of unemployment—is a production function with the number of jobseekers and the number of positions open for recruiting taken as inputs and the flow of newly matched worker-employer pairs as the output. Matching efficiency is a multiplicative shifter of the production function, analogous to the Hicks-neutral productivity index in production theory.

The term *mismatch* often appears in discussions of high unemployment. Shocks that cause widespread job loss and leave many workers unmatched with employers will generate mismatch. The role of the matching function is to cure mismatch by using resources—jobseekers' time and employers' recruiting expenditures. Thus mismatch is organic to labor-market models built on matching functions. The presence of high levels of unemployment is not necessarily a sign of a decline in matching efficiency. The appropriate way to proceed is to measure matching efficiency using standard ideas from production theory. If measured efficiency declines, a rising incidence of mismatch is one of a number of potential sources. Proper measurement of matching efficiency is a crucial starting point for understanding the sources of high unemployment.

The Beveridge curve is another way to characterize changes in matching efficiency. A decline in efficiency shifts the curve outward, so vacancies are higher for a given level of unemployment. We show how our results map into the Beveridge curve. The outward shift of the curve is the result of trends present during 2001 through 2007, not a special change in the crisis and post-crisis years, 2008 through 2012.

Most analysis of the U.S. labor market in the matching-function framework has taken unemployment to be the appropriate measure of jobseeking in the population. But it is well known that this view is incomplete. In the Current Population Survey (CPS) in 2006, the distribution of hires into new jobs was 21 percent from unemployment, 42 percent from people not previously in the labor force, and 37 percent from workers in previous jobs who

took new jobs without intervening unemployment or time out of the labor force. Job-to-job hiring has long been an important part of DMP modeling, but not in the measurement of matching efficiency. The remarkably large flow into jobs of people who were not previously counted as active searchers in the CPS has received less attention. An important exception is Veracierto (2011), a paper that we build on.

We develop the theory of aggregation of matching functions across diverse groups. The condition for aggregation is a natural one: changes in the success rates for job-seekers should move in proportion to one another. Our main finding is that matching efficiency measured consistently with our aggregation theory fell only slightly in recent years, and by no more than would have been expected from the earlier modest downward trend in efficiency. Earlier mis-measurement of matching efficiency was the result of treating jobseekers as homogeneous. Proper treatment of heterogeneity by reason for unemployment and duration of unemployment to date reverses the finding of a collapse of matching efficiency.

With the exception of Krueger, Cramer and Cho (2014), research on labor turnover has tended to focus on month-to-month changes in labor-market status—Blanchard and Diamond (1990) is a leading example. Because the separation rate from brand-new jobs is extremely high, the probability of employment a few months later conditional on unemployment in a given month is not as high as one might expect from the monthly job-finding rate. For example, the monthly job-finding rate for workers who recently suffered the loss of a permanent job was 34 percent in 2007. But measured over a three month span, only 47 percent of those workers held jobs at the end of the span. With average separation rates, 66 percent would have been holding jobs after two more chances of landing jobs with a probability of 34 percent. And 15 months later, with 12 additional chances at a 34 percent success rate, only 62 percent were holding jobs, against 85 percent with normal rates of losing or leaving jobs. Accordingly, we study job-finding rates over the full 15-month history of each worker in the CPS. We find that there has been an upward trend in matching efficiency measured by the longer-span measures of matching success (12 through 15 months after the conditioning date) compared with the shorter-span measures (one to three months after that date).

The appendix describes some of the many earlier papers on the topic of this paper.

1 Aggregating Matching Functions

A matching function is a function $m(X, V)$, increasing and weakly concave in the number of jobseekers X and the number of vacancies V . $H = m(X, V)$ is the flow of new hires emerging from the matching process. Most investigators take the function to have constant returns to scale. The job-seeking success hazard associated with m is

$$f = \phi\left(\frac{V}{X}\right) = \frac{m(X, V)}{X} = m\left(1, \frac{V}{X}\right). \quad (1)$$

f is the flow rate into new jobs of members of the homogeneous population measured by X .

Now we consider a heterogeneous set of jobseekers of various types. Type i has a matching efficiency parameter μ_i and a parameter ψ_i that indicates what fraction of the population P_i of type i are jobseekers. We define the effective number of jobseekers:

$$X = \sum_i \mu_i \psi_i P_i. \quad (2)$$

We assume that all the job-seekers search in the same market and have the same matching rate except for the efficiency parameter μ_i :

Assumption. Scaled matching hazard function and common pools of vacancies and competing jobseekers:

$$H_i = \mu_i \psi_i \phi\left(\frac{V}{X}\right) P_i. \quad (3)$$

Total hires are $H = \sum_i H_i$. Our basic result is:

Aggregation Theorem: Let m be the matching function corresponding to the jobseeking success hazard function ϕ . Then $H = m(X, V)$.

proof:

$$H = \sum_i H_i = \sum_i \mu_i \psi_i \phi\left(\frac{V}{X}\right) P_i = \phi\left(\frac{V}{X}\right) X = m(X, V). \quad (4)$$

We do not consider the distinction between a contact of a jobseeker and employer and the creation of a job match. The matching function takes account of the fact that many contacts do not result in hires.

Only the product of μ_i and ψ_i appears in these equations, not the two measures separately. There is no prospect of distinguishing changes in matching efficiency from changes in search propensities. From this point forward, we define γ_i as the product $\mu_i\psi_i$. We refer to γ_i as efficiency, but it should be kept in mind that a decline in our measure of efficiency may arise from a decline in the search propensity of a type rather than a decline in the efficiency of the search of those choosing to search.

1.1 Applying the aggregation principle

Petrongolo and Pissarides (2001) discuss the evidence that the matching function has the Cobb-Douglas form, where the elasticities with respect to X and V are η and $1 - \eta$:

$$H = X^\eta V^{1-\eta}. \quad (5)$$

The aggregate matching function has no efficiency parameter in our setup—efficiency shows up in the job-finding rates by type and is buried inside the aggregate effective count of jobseekers, X . We solve out X to get

$$\phi\left(\frac{V}{X}\right) = \left(\frac{V}{H}\right)^{\frac{1-\eta}{\eta}}, \quad (6)$$

which leads to

$$f_{i,t} = \gamma_{i,t} \left(\frac{V_t}{H_t}\right)^{\frac{1-\eta}{\eta}} = \gamma_{i,t} T_t, \quad (7)$$

where

$$T_t = \left(\frac{V_t}{H_t}\right)^{\frac{1-\eta}{\eta}}, \quad (8)$$

our measure of tightness. Finally,

$$\gamma_{i,t} = \frac{f_{i,t}}{T_t}. \quad (9)$$

We discuss the estimation of the elasticity η in a later section.

2 Job-Finding Rates

The standard concept of a job-finding rate is the probability that a job-seeker will find a job in a given month. We include rates based on that definition, but we also generalize it to study longer time spans, up to the longest found in the CPS. That span is 15 months, comparing the month the person entered the survey to the last month the person was in the survey.

<i>Calendar month</i>	<i>CPS month</i>	<i>Span, months</i>	<i>Unemployment duration, months</i>
November-08			0
December-08			1
January-09			2
February-09	1		3
March-09	2	0	4
April-09	3	1	5
May-09	4	2	6
June-09		3	7
July-09		4	8
August-09		5	9
September-09		6	10
October-09		7	11
November-09		8	12
December-09		9	13
January-10		10	14
February-10	13	11	15
March-10	14	12	16
April-10	15	13	17
May-10	16		18
June-10			19

Table 1: Example of CPS Survey Months, a Span, and an Unemployment Spell

We use the term *span* to mean the number of months between one observation on a person’s labor-market status and a subsequent observation. For example, the CPS might determine that a person was unemployed on account of the loss of a permanent job in March 2009 and unemployed as well in April 2010. The span in our sense would then be 13 months. It is important to understand that span is different from, for example, the duration of unemployment. In this example, the person might have been unemployed since November 2008 and thus had a duration of unemployment of four months as of March 2009 and 17 months as of April 2010. The beginning of a span is not necessarily in the month the person entered the CPS. In the example, the person could have entered the CPS in February 2009, so that the span began in the second month of the person’s period in the CPS and ended in the 15th month in the CPS. Table 1 shows the relation between the span, the CPS months, and the months of the spell of unemployment, in this example.

Over these spans, we focus on the experiences of people who were in a given labor-market status, such as looking for work after having recently quit a job. We define these statuses precisely in the next section. We then examine the probability that such a person would be employed, say, 12 months later. Longer spans matter for measuring job-finding success because many jobseekers find brief jobs, lasting only a few weeks or a month or two. A

job lasting a month counts as much as a job lasting years if the measure of success uses a one-month span. Longer spans give higher weight to longer-lasting jobs.

To see this, consider a simple model of labor-market turnover. There are two kinds of jobs, short and long. Jobseekers have a 30 percent monthly probability of taking a short job and a 10 percent probability of taking a long job. The monthly probability that a short job will end is 40 percent, and the probability that a long job will end is two percent. The mix of jobs held by workers one month after a time when they are looking for work but not working is three-fourths short and one-fourth long (the distribution across workers conditional on not working in the previous month and working this month). That fraction switches to one-third short and two-thirds long with a 12-month span, as can be calculated from the 12th power of the transition matrix of the Markov process defined by the transition probabilities.

In the formalization of our setup, the job-finding rate $f_{i,t,\tau,x}$ is the probability that a worker in status i in month t with personal characteristics x is employed in month $t + \tau$. We let this probability depend on a large vector of observed worker characteristics. The CPS sample is too small to estimate the probabilities nonparametrically, conditional on each possible combination of characteristics. Instead, we specify the probabilities as logit functions of the vector x , with time effects captured by time dummies. We allow different coefficients on the time dummies and worker characteristics for each origin status i and each time span τ . Thus, we assume

$$f_{i,t,\tau,x} = \frac{\exp(\kappa_{i,t,\tau} + x'\beta_{i,\tau})}{1 + \exp(\kappa_{i,t,\tau} + x'\beta_{i,\tau})}, \quad (10)$$

where $\kappa_{i,t,\tau}$ is the time effect at date t for workers in status i and a span of τ months. For job-to-job transitions, we define job-seeking success as being in a different job at the end of the span from the job at the beginning. With a one-month span, this definition is the same as the standard job-to-job rate. We can measure job-seeking success in the job-to-job case only over spans up to three months because the CPS does not keep track of respondents' employers during the eight-month gap between waves of interviews.

In a small number of cases where all respondents who started in status i in month t were employed at $t + \tau$ or where none of them were, we take the predicted job-finding rate to be 1 or 0.

A substantial literature describes reporting errors in the CPS and similar longitudinal surveys. Random errors in assigning workers to labor-market statuses result in overstatements of month-to-month transition rates. Correction of some of these errors is possible because of

redundancies in the data, but most escape detection except through re-interviews. A number of proposals have appeared in the literature to make corrections in population fractions based on heuristics, such as Abowd and Zellner (1985) and Poterba and Summers (1986). More recently, formal models of identified classification errors have appeared in the econometrics literature, such as Feng and Hu (2013). We do not find either of these approaches compelling. We do not think that any realistic model with classification errors is identified by longitudinal data alone. We believe that our approach based on studying longer-span conditional probabilities of employment solves at least part of the problem, in that transitory misclassification in the destination status will be unimportant for our longer-span measures. We do retain conditioning on a single-month measure of the origin status, which results in some blurring of our results.

3 Data

We use data from the monthly CPS for November 1999 through March 2014. These data permit the calculation of job-finding rates for individuals who started their searches in the years 2001 through 2012.

Because the CPS interviews households for four consecutive months, skips the next 8 months, then interviews again for four months, each person covered for every scheduled interview contributes 6 observations spanning single months, 4 spanning two months, 4 spanning 12 months, and one spanning 15 months, to give a few examples. In principle, we can study job-seeking spans of one, two, three, nine, 10, 11, 12, 13, 14, and 15 months. For simplicity, we omit the nine-, 10- and 11-month spans and focus on the short spans from one through three months and the long spans from 12 through 15 months.

The CPS divides the civilian noninstitutional population, ages 16 and older, into people who are employed, unemployed, and not in the labor force. Employed people are those who worked for pay or profit during the reference week, were temporarily absent from work for reasons such as vacation, illness, weather, or industrial dispute, or did at least 15 hours of unpaid work in a family-owned business. People who are not employed are classified as unemployed if they are currently available for work and either have actively looked for work during the previous four weeks or expect to be recalled from a temporary layoff. All other people who are not employed are classified as not in the labor force. We further divide the

unemployed people according to the reasons they became unemployed and the length of time since that happened. We derive a total of 15 labor-market statuses. The first two are:

- *Out of labor force*: people who did not satisfy the CPS definition of either employed or unemployed.
- *Working*: employed people.

The next set of statuses is for people who have been unemployed for three weeks or less:

- *Recently laid off*: unemployed people who have been on furlough for three weeks or less from an earlier job, with the possibility of recall.
- *Recently lost permanent job*: people who lost jobs within the previous three weeks, not on layoff or separated from a temporary job, who were working or left military service immediately before they began looking for work.
- *Temp job recently ended*: unemployed people, not on layoff, whose last jobs were explicitly temporary and ended within the past three weeks or less.
- *Recently quit*: unemployed people who quit their last jobs within the past three weeks.
- *Recently entered*: unemployed people who have never worked and who started looking for work within the past three weeks.
- *Recently re-entered*: unemployed people, who started looking for work within the past three weeks, who were not working or in military service immediately before they began looking for work, but who have worked at some time in the past.

The following categories parallel those above, with duration of unemployment to date of 4 to 26 weeks:

- *On layoff for months*
- *Lost permanent job months ago*
- *Temp job ended months ago*
- *Quit months ago*
- *Entered months ago*

- *Re-entered months ago*

The last category is

- *Long-term unemployed*: those unemployed to date more than 26 weeks.

We do not separate the long-term unemployed by reason for unemployment because, at most times, the number of long-term-unemployed respondents in the CPS is too small to estimate probabilities reliably if we further disaggregate those respondents by reason for unemployment.

We match respondents across months using the method of Nekarda (2009). Nekarda's approach considers the full set of eight monthly observations that potentially come from the same person and assigns to each observation a probability of actually coming from the same person, based on the recorded information on the person's race, sex, and age. This probability, combined with the survey weights, is used to weight the observed transitions when we compute job-finding rates. Relative to methods such as that of Madrian and Lefgren (2000), which label respondents as matched or not across each consecutive pair of months, Nekarda's method is more suitable for measuring job-finding rates across long time spans because errors in recording race, sex, and age during intervening months are less likely to break the match.

We remove high-frequency, likely spurious transitions between unemployment and non-participation following Elsbj, Hobijn and Şahin (2013). Specifically, if a respondent is out of the labor force, unemployed, and out of the labor force in three consecutive months, we recode the middle month to *out of the labor force*. If the respondent is unemployed in the first and third months and out of the labor force in the middle month, we recode the middle month to unemployed with the same reason for unemployment as the first month. Among respondents who remain unemployed, we remove spurious changes in the reason for unemployment by requiring that the reason must remain the same as that given in the first interview of the unemployment spell, except that we allow transitions between temporary layoff status and permanent job loss after one month of unemployment because a worker could be temporarily laid off and later learn that the job loss had become permanent. We do not allow transitions between temporary layoff and permanent job loss once unemployment duration exceeds one month because too few such transitions are in the raw data to allow us to estimate the logit model if we allow them.

The CPS allows workers who enter unemployment to report a positive initial duration. Elsby, Hobijn, Şahin and Valletta (2011) show that inflows to high-duration unemployment are essential to understanding labor market flows during the Great Recession. We therefore accept those observations. This procedure implies that unemployment duration should not be interpreted literally as duration of the current spell, but rather an indicator of the time that has elapsed since the individual has held a job more durable than an interim job.

The variables describing personal characteristics, denoted $x_{k,t}$, are dummy variables for

- female
- married
- six age groups—16–24, 25–34, 35–44, 45–54, 55–64, and 65-plus
- four education groups—less than high school, high school graduate, some college but less than a bachelor’s degree, and bachelor’s or higher degree
- five unemployment duration groups, for the equations describing job-finding conditioned on unemployment of 4 to 26 weeks—categories are 4–8 weeks, 9–13 weeks, 14–17 weeks, 18–21 weeks, and 22–26 weeks

We compute approximate bootstrap standard errors for our estimates. We recompute all of the estimates in 100 bootstrap samples, which we construct as follows: Define a state-month as the set of all households in a given state of the U.S. whose first interview fell in a given month. We create the bootstrap samples by resampling households with replacement within each state-month. Each resampling follows the individual through all subsequent appearances in the CPS. This procedure accounts for the stratification of the CPS sample by state. It amounts to a block-bootstrap within households and thus accounts for the correlations across members and over time within each household. It also accounts for our use of overlapping transitions—for example, our estimates of the two-month job-finding rate uses transitions from the first to third month and from the second to the fourth month for the same person. Following Rao, Wu and Yue (1992), we resample $n_h - 1$ households from a state-month with n_h households in the original sample so that the bootstrap is unbiased. We use Kolenikov’s (2010) Stata program to construct the bootstrap samples. Because we do not have access to some of the underlying data that the Census Bureau uses to construct poststratified survey weights in the CPS, our bootstrap samples cannot account for the

impact of the poststratification procedure. This omission is likely to inflate our bootstrap standard errors because the poststratification procedure reduces variance by holding constant the distributions of some demographic variables.

The rare event of a sample size of zero within a status-month-span cell occurred once in the CPS data. No individuals who are new entrants to the labor force in February 2008 were present for a full 15-month time span. As a result, we cannot estimate the time effect in $\kappa_{i,t,\tau}$ in equation (10) for that initial status, date, and time span. Instead, we impute the 15-month job-finding rates for new entrants in February 2008 based on the job-finding rates in adjacent months and years. Specifically, we impute

$$f_{i,\text{Feb } 2008,15} = \frac{1}{2} \left(\frac{f_{i,\text{Feb } 2007,15}}{f_{i,\text{Jan } 2007,15} + f_{i,\text{Mar } 2007,15}} + \frac{f_{i,\text{Feb } 2009,15}}{f_{i,\text{Jan } 2009,15} + f_{i,\text{Mar } 2009,15}} \right) (f_{i,\text{Jan } 2008,15} + f_{i,\text{Mar } 2008,15}),$$

where $i = \textit{recently entered labor force}$. We apply a similar procedure in the bootstrapped job-finding rates when a particular bootstrap sample has no observations for a given initial status, date, and time span.

4 Estimated Job-Finding Rates

Our estimation yields a great mass of logit coefficients, available from the online backup for the paper. In this section, we display and interpret the results in terms of calculated job-finding rates adjusted for changing composition of the labor force. We make the adjustment by choosing a base period, January 2005 to December 2007. We calculate the distribution of personal characteristics x across all respondents in the base period. Then, for each month from 2001 through 2012, we calculate the fitted job-finding probabilities from the logits separately for each possible vector of personal characteristics. Finally, we compute the average probabilities across the distribution of personal characteristics measured in the base period.

Figure 1 shows the mix-adjusted estimated job-finding probabilities for one important initial status, *recently lost permanent job*, along with the bootstrap standard errors for these probabilities. The lowest curve is the probability that a person who lost a permanent job in the past six months and has been searching since then will be employed one month later. The probability runs around 30 percent. It fell in the recession of 2001, rose to a peak in 2005, fell again in the Great Recession, and rose only a bit in the recovery through 2012. The probability has a noticeable downward trend.




Figure 1: Employment Probabilities Subsequent to Short-Duration Unemployment from Loss of a Permanent Job

Source: Authors' calculations from Current Population Survey microdata. Annual averages of monthly data. Vertical bars show a range of plus or minus one standard error around point estimates.

The next curve up is the probability that a person will be re-employed after two months. The curve is close to parallel with the one-month curve, and only slightly above the one-month curve. In 2007, the one-month probability was 33.6 percent and the two-month probability was 43.0 percent. If the monthly job-finding rate was truly 33.6 percent and if there was no chance of losing a job in the second month that had been found in the first month, the probability of being employed in the second month would be $0.336 + (1 - 0.336) \times 0.336 = 0.559$, far above the actual value.

As far as we know, Krueger et al. (2014) were the first investigators to note this anomaly. They studied long-term unemployment. They concluded, "...the long-term unemployed face difficulty regaining full-time, steady work over the longest period we can observe in CPS data. It appears that reemployment does not fully reset the clock for the long-term unemployed." Our results show that the same proposition applies to every type of unemployment.

The remaining curves in Figure 1 lie even closer to each other, so the anomaly is even more acute for longer spans. One reason that the multi-month probabilities are so far below their hypothetical levels may be misclassification in the CPS. Errors could take two forms.

One is classifying people as unemployed when they are actually employed. Though this type of error would exaggerate one-month employment probabilities, on the assumption that the error would have a probability of correction in the next month, the exaggeration would apply to longer spans as well. For example, suppose that these misclassifications are corrected in the succeeding month and suppose that the jobs have close to zero separation rates. Then, following a misclassification, a long series of observations of employment would occur. There would be an equal upward bias for all of the employment probabilities. So misclassification of the initial status of respondents is not a likely explanation for the anomaly.

The second type of error misclassifies job-seekers as employed when they are actually still unemployed, in months after the initial conditioning month. If such errors are prevalent and transitory, the anomaly would be explained. High measured job-finding rates based on month-to-month changes would be an illusion of phantasmal jobs, so brief that they would not show up in the longer-span probabilities.

The other explanation—the one we tend to favor—is that the brief jobs recorded in the CPS are true jobs, but truly brief. Hall (1995) proposed that brief interim jobs were part of the experience of the unemployed. Hall (2014a) shows that the incidence of very short jobs among newly filled jobs is remarkably high, based on the number of respondents in the CPS who report short job tenure. Hyatt and Spletzer (2013) provide evidence from a variety of sources on the decline in short-duration jobs. Davis and Haltiwanger (2014) report lower turnover rates of labor without commenting on the role of short-duration jobs in that trend, and condemn the trend as a decline in the fluidity of the labor market. We question that interpretation, because our results show higher success rates for search for higher duration jobs and point in the direction that the trend arises from declining rates of very short-term interim jobs in the re-employment process.

Table 2 and Table 3 show the estimated employment success rates for the year 2007 by initial status. The probabilities are computed separately for each month of the year and averaged over the 12 months. For each status, the row labeled Actual gives the percent of a random sample of people in that status in a given month who are employed in the later months of the CPS schedule. For example, 4.7 percent of those out of the labor force in a given month are employed in the following month and 13.0 percent 15 months later. The row labeled Benchmark is the projected percentage if the job-finding rate for month 1 applies in all the later months, but there is a monthly probability of 6 percent that any job found

<i>Initial status</i>		<i>Percent unemployed as of a later month</i>						
		<i>Months later</i>						
		1	2	3	12	13	14	15
Out of labor force	Actual	4.7	6.3	7.4	12.0	12.4	12.8	13.0
	(Standard error)	(0.0)	(0.1)	(0.1)	(0.1)	(0.2)	(0.2)	(0.2)
	Benchmark	4.7	9.0	12.7	32.8	34.0	35.1	36.1
Recently laid off	Actual	56.0	64.9	64.9	62.2	60.3	58.4	61.8
	(Standard error)	(1.5)	(1.7)	(2.3)	(1.8)	(2.1)	(2.4)	(3.3)
	Benchmark	56.0	77.3	85.4	90.3	90.3	90.3	90.3
Recently lost permanent job	Actual	33.6	43.0	46.6	66.4	62.5	59.7	61.5
	(Standard error)	(1.5)	(2.0)	(2.4)	(2.2)	(2.4)	(3.2)	(4.2)
	Benchmark	33.6	53.9	66.2	84.7	84.7	84.8	84.8
Temp job recently ended	Actual	42.0	54.2	49.3	59.7	61.4	66.4	58.0
	(Standard error)	(2.3)	(2.8)	(4.2)	(2.7)	(3.9)	(4.5)	(5.6)
	Benchmark	42.0	63.9	75.2	87.5	87.5	87.5	87.5
Recently quit a job	Actual	40.4	51.7	58.1	69.1	64.1	67.6	58.7
	(Standard error)	(2.2)	(2.6)	(3.2)	(2.5)	(3.2)	(4.0)	(4.7)
	Benchmark	40.4	62.1	73.7	87.0	87.0	87.1	87.1
Recently entered LF	Actual	29.3	28.9	25.6	37.0	41.3	37.5	43.0
	(Standard error)	(2.8)	(2.9)	(3.5)	(4.0)	(4.8)	(4.8)	(7.1)
	Benchmark	29.3	48.3	60.6	82.6	82.7	82.8	82.9
Recently re-entered LF	Actual	35.5	44.0	43.7	52.5	56.0	56.6	57.0
	(Standard error)	(1.3)	(1.4)	(2.1)	(2.2)	(2.8)	(2.7)	(3.4)
	Benchmark	35.5	56.3	68.4	85.4	85.5	85.5	85.5

Table 2: Subsequent Employment Probabilities by Initial Status, Actual and Benchmark, 2007: Out of Labor Force and Recently Unemployed

		<i>Percent unemployed as of a later month</i>						
<i>Initial status</i>		<i>Months later</i>						
		1	2	3	12	13	14	15
On layoff for months	Actual	42.5	51.2	59.2	49.8	54.4	63.0	63.8
	(Standard error)	(1.6)	(2.0)	(2.5)	(2.2)	(2.3)	(2.9)	(3.8)
	Benchmark	42.5	64.3	75.6	87.6	87.6	87.6	87.6
Lost permanent job months ago	Actual	22.6	31.3	37.8	58.7	58.7	59.1	56.3
	(Standard error)	(0.8)	(1.1)	(1.7)	(1.7)	(1.8)	(1.9)	(2.5)
	Benchmark	22.6	38.8	50.3	77.7	78.1	78.4	78.6
Temp job ended months ago	Actual	26.4	33.0	37.4	49.3	50.1	50.8	44.8
	(Standard error)	(1.3)	(2.1)	(2.6)	(2.4)	(2.6)	(3.1)	(4.3)
	Benchmark	26.4	44.3	56.3	80.8	81.0	81.2	81.3
Quit a job months ago	Actual	27.0	35.4	42.5	65.1	64.9	63.2	65.6
	(Standard error)	(1.2)	(1.7)	(2.4)	(2.1)	(2.4)	(2.8)	(3.2)
	Benchmark	27.0	45.1	57.2	81.2	81.4	81.5	81.6
Entered LF months ago	Actual	16.9	21.3	27.8	40.8	45.0	41.7	38.5
	(Standard error)	(1.3)	(1.9)	(2.4)	(2.9)	(3.3)	(3.9)	(4.4)
	Benchmark	16.9	29.9	39.9	70.5	71.2	71.8	72.3
Re-entered LF months ago	Actual	24.0	31.5	35.8	50.2	51.1	51.0	48.9
	(Standard error)	(0.7)	(0.9)	(1.3)	(1.8)	(1.8)	(2.0)	(2.7)
	Benchmark	24.0	40.8	52.6	78.9	79.2	79.5	79.6
Long-term unemployed	Actual	15.6	22.0	25.9	35.6	36.7	37.3	34.6
	(Standard error)	(0.7)	(0.9)	(1.3)	(1.6)	(1.9)	(2.0)	(2.2)
	Benchmark	15.6	27.9	37.5	68.4	69.2	69.9	70.4

Table 3: Subsequent Employment Probabilities by Initial Status, Actual and Benchmark, 2007: Unemployed for Months and Long-Term

ends in a subsequent month and the worker cycles back to the status named at the left. Six percent per month is the typical job separation rate found in the CPS. For all initial cases and all spans of 2 months or more, the actual employment rate falls short of the benchmark, often by large amounts. For example, for workers starting in the *recently laid off* status, which has a high one-month job-finding rate of 56.0 percent, the benchmark would have 90.3 percent back at work 15 months later, but in fact, only 61.8 percent are back. The separation rates needed to explain the observed employment probabilities are in the range of 50 or even 70 percent per month.

Table 4 summarizes our findings for employment probabilities conditional on originating in each of the job-seeking statuses. The left panel shows the probabilities averaged over the early three months following the conditioning month and the right panel over the later four months. The third column in each panel shows the ratio of the employment probability in 2012 to the probability in 2001—these ratios are good measures of the trend because the estimated annual probabilities lie along smooth lines. In almost all originating statuses, the trend is downward in the probabilities measured up to 3 months after the conditioning month; the one exception is the originating status *recently laid off*, for which the trend is flat. By contrast, the probabilities measured 12 to 15 months after the conditioning month, in the right-hand panel, generally have smaller downward trends and in some cases upward trends. Success rates in finding first jobs following spells of job search have declined over time, while success rates for finding jobs over longer periods of search have risen. As we noted earlier, longer-span employment probabilities are better at capturing success in finding longer-duration jobs.

The employment probabilities in Table 4 vary over a wide range across the conditioning statuses. Not including the employed, for whom we look at the probability of changing jobs, the lowest job-finding rate is for people starting in the status *out of the labor force*. In 2012, their short-span subsequent employment probability was 5.0 percent and their long-span rate was 10.7 percent. Most people classified as out of the labor force remain in non-market activities from one year to the next. The CPS inquires about job-seeking interest among these people, and subsequent employment probabilities are higher among those indicating interest, but we do not pursue that topic in this paper. It would be important for any attempt to place the measurement of unemployment on the footing proposed in Flinn and Heckman (1983).

<i>Initial status</i>	<i>Average employment probability, months 1 to 3</i>			<i>Average employment probability, months 12 to 15</i>		
	2001	2012	<i>Ratio</i>	2001	2012	<i>Ratio</i>
Out of labor force (Standard error)	6.9 (0.1)	5.0 (0.1)	0.73 (0.01)	13.3 (0.2)	10.7 (0.2)	0.80 (0.01)
Recently laid off (Standard error)	59.7 (1.7)	60.5 (1.8)	1.01 (0.03)	64.1 (2.6)	67.8 (2.6)	1.06 (0.05)
Recently lost permanent job (Standard error)	39.4 (2.0)	36.1 (2.0)	0.92 (0.05)	65.7 (3.2)	61.5 (3.2)	0.94 (0.05)
Temp job recently ended (Standard error)	49.6 (3.2)	43.2 (2.9)	0.87 (0.07)	61.5 (4.3)	57.9 (3.8)	0.94 (0.06)
Recently quit a job (Standard error)	55.0 (2.1)	46.0 (3.0)	0.84 (0.05)	68.6 (3.7)	65.4 (4.6)	0.95 (0.07)
Recently entered LF (Standard error)	30.8 (3.9)	15.2 (2.1)	0.49 (0.07)	36.6 (5.9)	38.0 (5.0)	1.04 (0.18)
Recently re-entered LF (Standard error)	41.9 (1.6)	31.0 (1.9)	0.74 (0.04)	53.5 (2.6)	52.7 (3.3)	0.99 (0.06)
On layoff for months (Standard error)	50.8 (1.7)	46.6 (1.6)	0.92 (0.04)	63.9 (2.4)	57.6 (2.7)	0.90 (0.05)
Lost permanent job months ago (Standard error)	30.3 (1.0)	23.2 (0.8)	0.76 (0.03)	63.9 (1.9)	55.6 (1.6)	0.87 (0.03)
Temp job ended months ago (Standard error)	34.6 (2.2)	31.4 (1.7)	0.91 (0.06)	55.5 (3.7)	49.2 (2.7)	0.89 (0.07)
Quit a job months ago (Standard error)	37.2 (2.0)	31.1 (1.8)	0.84 (0.06)	59.1 (3.1)	52.6 (3.3)	0.89 (0.06)
Entered LF months ago (Standard error)	20.0 (2.3)	15.4 (1.4)	0.77 (0.10)	39.3 (4.4)	39.9 (3.0)	1.01 (0.12)
Re-entered LF months ago (Standard error)	31.8 (1.2)	24.5 (1.0)	0.77 (0.04)	51.3 (1.9)	43.0 (1.9)	0.84 (0.04)
Long-term unemployed (Standard error)	23.0 (1.4)	15.2 (0.5)	0.66 (0.04)	39.4 (2.4)	35.1 (1.0)	0.89 (0.06)
Employed (probability of changing job) (Standard error)	5.8 (0.1)	4.5 (0.1)	0.77 (0.01)			

Table 4: Subsequent Employment Probabilities for Short and Long Spans, 2001 and 2013, with Growth Ratio

The long-term unemployed had short-span re-employment success rates of only 15.2 percent in 2012. Over the longer span of 12 to 15 months after the conditioning month (which is itself at least 6 months after the job loss), 35.1 percent of the this group was employed. Though these figures make it clear that workers who fail to find jobs after six months of unemployment are not very likely to find jobs after another year of search, that proposition was true in all earlier years as well, including 2001, a year of substantially lower overall unemployment than 2012.

Entrants and re-entrants tend to have lower employment probabilities than other categories of unemployment apart from long-term unemployment. Those who lost permanent jobs, either recently or months ago, have quite low short-span success rates but longer-span rates comparable to other categories of unemployed job-seekers.

5 Job-Finding Rates and Tightness

5.1 Basic equation for estimation of the elasticity of the job-finding rate with respect to tightness

Equation (7) leads to the following model of the measured log job-finding rate over a τ -month span:

$$\log f_{i,t,\tau} = \log \gamma_{i,t,\tau} + \nu_\tau \log d_t + \epsilon_{i,t,\tau}^m, \quad (11)$$

where $\epsilon_{i,t,\tau}^m$ is a measurement error. Here ν_τ is the elasticity of job-finding with respect to the measure of tightness from the employer's perspective, $d_t = V_t/H_t$, which is the duration of vacancies, measured as the ratio of the stock of vacancies to the flow of hires. The elasticity is related to the elasticity of the matching function as $\nu = (1 - \eta)/\eta$.

We assume that matching efficiency satisfies

$$\log \gamma_{i,t,\tau} = \alpha_{i,\tau} + \delta_{i,\tau}t + \psi_{i,s} + \xi_{i,t,\tau}, \quad (12)$$

where s is the month of the year, to capture seasonal effects, $\psi_{i,s}$, and t is time measured in months, to capture a trend, $\delta_{i,\tau}t$. The model we estimate is thus

$$\log f_{i,t,\tau} = \alpha_{i,\tau} + \delta_{i,\tau}t + \psi_{i,s} + \nu_\tau \log d_t + \epsilon_{i,t,\tau}, \quad (13)$$

where

$$\epsilon_{i,t,\tau} = \epsilon_{i,t,\tau}^m + \xi_{i,t,\tau}. \quad (14)$$

5.2 Identification

Our first identifying assumption is

$$\mathbb{E}(\epsilon_{i,t,\tau}|t) = 0, \tag{15}$$

so the month, t , is eligible as an instrumental variable and seasonal dummies are also eligible as instruments.

The job-finding rate and labor-market tightness are obviously jointly determined, so a further assumption about the disturbance $\epsilon_{i,t,\tau}$ is required for identification—the disturbance is not plausibly orthogonal to either variable. Our second identifying assumption is that $\epsilon_{i,t,\tau}$ is orthogonal to the log of real GDP. This assumption is likely to hold at least for one major source of correlation between the disturbance and the variables, namely measurement error. We use the monthly estimate of real GDP from Macroeconomic Advisers.

5.3 Further aspects of estimation

We average the three short spans (one, two, and three months after the conditioning status) to form the job-finding rate for the first span category, called *short*, and the four longer spans (12 through 15 months) to form the second job-finding rate category, called *long*. For the short job-finding rate, we can include in our data the job-changing rate for those starting in the *employed* status. For the long job-finding rate, we cannot calculate the job-changing rate; thus, for comparability between the short and long equations, we also estimate the short equation without including the job-changing rate. We estimate equation (13) with the instrumental variables noted above, using monthly data on job-finding rates.

As equation (13) indicates, we pool the data for initial statuses in estimation, to enforce the implication of the model that the elasticity of the job-finding rate with respect to vacancy duration, ν_τ , is the same across those statuses, though different between the short and long spans. We do not take into account any correlation of the disturbances across the statuses. Thus our estimates are unbiased but not minimum variance, if correlation is present. Because we use a bootstrap strategy to calculate standard errors that preserves the correlation, those standard errors are not biased by the correlation. The correlation is positive in almost all cases, but relatively mild—the average absolute values of the off-diagonal elements of the correlation matrices are 0.26 for short spans, without job-to-job, 0.27 with job-to-job, and 0.15 for long spans. We do not believe that a three-stage least squares estimation procedure

would be appropriate, given the fairly small sample size of 84 observations and 13 separate values of the seasonal dummies and the time trend in each equation.

The residuals from equation (13) form an index of detrended matching efficiency:

$$\epsilon_{i,t,\tau} = \log f_{i,t,\tau} - \alpha_{i,\tau} - \delta_{i,\tau}t - \psi_{i,s} - \nu_{\tau} \log d_t, \quad (16)$$

as the observed job-finding rate measured around its status- and span-specific constant level and trend, and adjusted for changes in labor-market tightness. These residuals also include measurement error in job-finding rates, but such measurement errors should average to zero over time.

We use the estimates of job-finding rates adjusted for the changing characteristics of the population, as discussed earlier, as the left-hand variable of equation (13). Although, in principle, it would be possible to combine the two estimation stages, we doubt its practicality and have no reason to believe it would affect our conclusions.

5.4 Measuring tightness, d

Figure 2 shows the number of new hires from the CPS and from the Job Openings and Labor Turnover Survey of employers, JOLTS. The CPS and JOLTS surveys vary similarly over time, but the level of hires is substantially higher in the CPS. The reasons for the discrepancy may include: (1) JOLTS does not include hires at new establishments or self-employment, as Davis, Faberman, Haltiwanger and Rucker (2010) discuss, and (2) the CPS may capture more of the hiring into jobs that last only days or a few weeks. Hires track the business cycle, but with fairly low amplitude.

Figure 3 shows the number of job openings (vacancies) from JOLTS. This series traces the business cycle with high amplitude—vacancies are high in tight market around peaks and low in slack markets around business-cycle troughs.

Figure 4 shows labor-market tightness, d , using the JOLTS measures of hires and vacancies. We do not use the CPS measure of hires to construct tightness because the CPS survey covers a larger universe of jobs than the JOLTS sample that we use for vacancies. Because vacancies vary more in proportional terms than do hires, the vacancy/hires ratio is quite procyclical.




Figure 2: Number of Monthly Hires, in Thousands, from the JOLTS and the CPS

5.5 Estimates

Table 5 shows estimates of the elasticity of the matching function based on equation (13). Because of the concern that the financial crisis and ensuing deep slump may have had special effects on matching efficiency, we estimate equation (13) for the period 2001 through 2007 and use it to calculate residuals for the entire period, 2001 through 2012. Below, we test for a possible parallel decline in matching efficiency during the recession and slump by estimating an equation that includes the later years. For the short-span equation, the estimate of 1.16, corresponding to a matching elasticity of 0.46, is in line with the estimates surveyed in Petrongolo and Pissarides (2001). We are not aware of any previous research on the longer-span matching-function elasticity.

5.6 Implied matching efficiency

We calculate indexes of matching efficiency for each of the 15 initial labor-market statuses. Because we hold the distribution of individuals' characteristics constant in calculating the job-finding rates on the left-hand side of equation (13), the movements in these indexes are insulated from changes in the distribution of characteristics. Figure 5 shows the resulting detrended indexes. These are the exponentials of the values described in equation (16) and




Figure 3: Number of Job Openings, in Thousands, from JOLTS




Figure 4: Labor-Market Tightness, Calculated from JOLTS




Figure 5: Detrended Matching Efficiency, 2001 through 2012, for Selected Initial Statuses

Source: Authors' calculations from Current Population Survey microdata. Annual averages of monthly data. Vertical bars show a range of plus or minus one standard error around point estimates. Measures for 1- to 3-month spans derived from estimates of equation (13) including data including job-to-job transitions.

<i>Monthly span of job-finding rate</i>	<i>Include job-to-job movers?</i>	<i>Elasticity with respect to vacancy duration</i>	<i>Implied elasticity of the matching function</i>	<i>Standard error of residuals</i>
Short	Yes	1.158 (0.206)	0.463 (0.048)	0.187 (0.010)
Short	No	1.164 (0.219)	0.462 (0.051)	0.191 (0.010)
Long	No	0.411 (0.237)	0.709 (0.134)	0.189 (0.016)

Table 5: Estimated Elasticities of Job-Finding with Respect to Market Tightness.

Source: Authors' calculations from Current Population Survey monthly microdata. Bootstrap standard errors in parentheses.

are indexes normalized to equal one in 2007. Recall that the trends are estimated through 2007 so they exclude the effects of the recession. The short-span results are derived from estimates of equation (13) that include data on job-to-job transitions.

The pattern of annual matching efficiency for the initial status *recently lost permanent job* is representative in terms of its movement over time and more precisely estimated because large numbers of jobseekers fell into this category. In that category, both measures of detrended efficiency rose during the recovery from the 2001 recession, fell as the economy reached its peak in 2007 (where the index is one by construction). After the recession matching efficiency as measured over short spans fell, while efficiency over long spans rose, though neither change was very large. We noted earlier that the measure over long spans gives more weights to longer-lasting new jobs, so the finding of improved efficiency for that measure suggests an improvement in labor-market performance that is not apparent in the conventional approach based on one-month spans. In the category *lost permanent job months ago*, the same pattern of declining efficiency over short spans and improving efficiency over long spans is present but larger. And in the closely watched category *long-term unemployed*, the same pattern is even stronger. By contrast, in the category *quit a job months ago*, the pattern is reversed—matching efficiency plunged by the long-span measure but rose a little by the short-span method.

Figure 6 shows the indexes without subtraction of the trend terms in equation (13). Notice that the trends are downward over time for essentially all of the initial statuses,

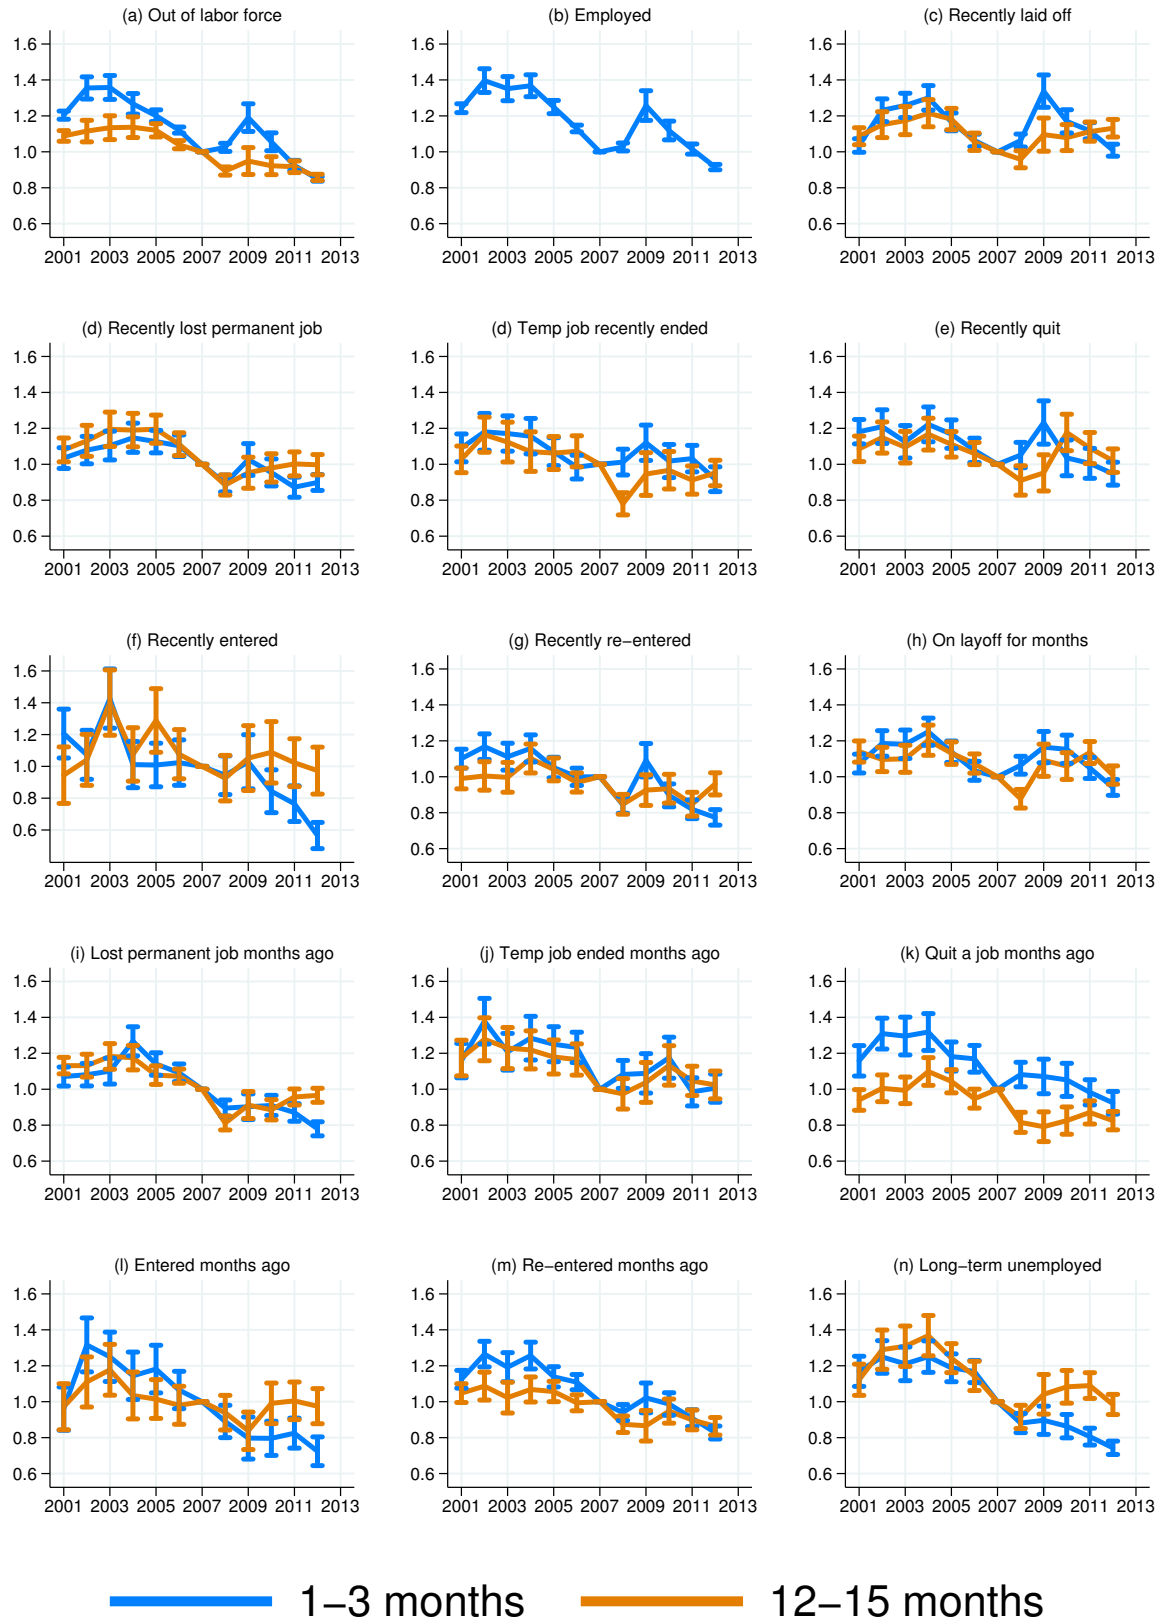


Figure 6: Non-detrended Matching Efficiency, 2001 through 2012, for Selected Initial Statuses

Source: Authors' calculations from Current Population Survey microdata. Annual averages of monthly data. Vertical bars show a range of plus or minus one standard error around point estimates. Measures for 1- to 3-month spans derived from estimates of equation (13) including data including job-to-job transitions.




Figure 7: Overall Matching Efficiency, 2001 through 2012

Source: Authors' calculations from Current Population Survey microdata. Annual averages of monthly data. Vertical bars show a range of plus or minus one standard error around point estimates.

corresponding to the ratios of 2012 job-finding rates to 2001 rates in Table 4 that are almost all below one.

Figure 7 shows indexes of matching efficiency across all of the initial statuses. The overall detrended index is a weighted average of the 15 detrended components in Figure 5, using weights representing the relative shares of the components in the population in the three years preceding the crisis, 2005 through 2007. Because the job-finding rates hold constant the distribution of worker characteristics conditional on labor market status, this aggregate index holds constant the joint distribution of worker characteristics and labor market status. The movements in matching efficiency measured by the aggregate index result from changes in the efficiency of particular types of workers, not in the distribution of workers. We construct three versions of the aggregate detrended matching efficiency index: an index for short spans that includes job-to-job movers, an index for short spans that includes only the unemployed and people not in the labor force, and an index for long spans that includes only the unemployed and people not in the labor force. Including job-to-job movers in the overall measure has little effect on the overall measure of matching efficiency because job-to-job movers' matching efficiency moved similarly to that of other people, as shown in Figure 5. The estimated indexes show that detrended matching efficiency for both measures based on short spans is quite cyclical, rising soon after the onset of recessions and then falling during recoveries. With adjustment for trend, short-span efficiency was essentially the same in 2012

	(1)	(2)	(3)
Elasticity with respect to vacancy duration	1.310 (0.213)	1.351 (0.227)	0.548 (0.256)
Implied elasticity of matching function	0.433 (0.042)	0.425 (0.044)	0.646 (0.125)
Standard error of residuals	0.194 (0.010)	0.199 (0.011)	0.192 (0.012)
Post-crisis year:			
2008	-0.059 (0.020)	-0.059 (0.022)	-0.168 (0.027)
2009	0.091 (0.074)	0.092 (0.079)	-0.035 (0.083)
2010	0.030 (0.059)	0.029 (0.063)	0.012 (0.066)
2011	-0.027 (0.046)	-0.030 (0.049)	0.001 (0.050)
2012	-0.109 (0.038)	-0.115 (0.041)	-0.026 (0.041)
Monthly span of job-finding rate	Short	Short	Long
Include job-to-job movers	Yes	No	No

Table 6: Year Effects in Matching Efficiency

Source: Authors' calculations from Current Population Survey microdata. Estimated on monthly data, with month-of-year dummies included in regression. Standard errors in parentheses.

as in 2007 and 2001. Long-span efficiency is less volatile, but was somewhat below its 2007 level in 2012, adjusted for pre-crisis trends.

The right-hand graph in Figure 7 shows the same data without adjustment for trend. Matching efficiency at both short and long spans has trended downward since 2001. This trend is more pronounced in the short-span measure.

5.7 Measuring the post-crisis movements of matching efficiency

A second approach to quantifying the post-crisis decline in matching efficiency and to test the hypothesis of no such effects is to estimate equation (13) on data from the full range of years, 2001 to 2012, and to include year effects for 2008 through 2012. Table 6 shows the results of this approach for the two monthly spans of job-finding rates. The top panel shows the estimated values of the elasticity of job-finding rates with respect to labor market tightness, ν_τ . The elasticity is somewhat higher in this estimate than in the original one, though the

differences are within the confidence intervals for the point estimates. The implied elasticity of the matching function is little changed, however. For both short and long spans, there were statistically unambiguous shortfalls of matching efficiency relative to the pre-crisis trends in 2008. This decrease in matching efficiency was about six percent for short-span job-finding rates and 17 percent for longer spans. For short spans, matching efficiency also fell about 11 percent below pre-crisis trends in 2012. In 2009 through 2011, the divergence from pre-crisis trends was statistically ambiguous.

6 Shifts in the Beveridge Curve

Many observers use the Beveridge curve, with unemployment on the horizontal axis and the vacancy rate on the vertical axis, to study the matching process. Our approach is to use the efficiency of the matching function as the more fundamental concept. In this section, we restate our findings in terms of shifts in the Beveridge curve.

Figure 8 shows the Beveridge curve based on the JOLTS measure of the vacancy rate (vacancies divided by the labor force) and the standard unemployment rate, for the years 2007 through 2012. For comparability with the rest of our analysis, we recalculate the unemployment rate in the CPS micro data; it differs by one-tenth of one percentage point from the published rate in some years because of our data cleaning steps. The figure displays the outward shift that has attracted so much attention. As the labor market first slackened and then tightened during the crisis and its aftermath, the vacancy rate and unemployment rate did not return to their starting points. Rather, there was a movement toward more unemployment given the vacancy rate, or, equivalently, higher vacancies given the unemployment rate. In 2012, the vacancy rate was about the same as in 2008, a year affected only toward the end by the crisis. But the unemployment rate was 8 percent in 2012, compared with less than 6 percent in 2008. The matching process was not working nearly as well in 2012 as in 2008, according to the Beveridge curve.

Figure 9 shows Beveridge curves constructed from the results in this paper. The horizontal axis is a comprehensive index of jobseeking, covering all the statuses in our model—it captures the three-quarters of successful job seeking carried out by people who are not counted as unemployed. The weights in the index are the exponentiated fixed-effect coefficients in our model for the 15 statuses, measured in 2007. As with our measure of aggregate matching efficiency, we construct three versions of the jobseeking index: one for short spans

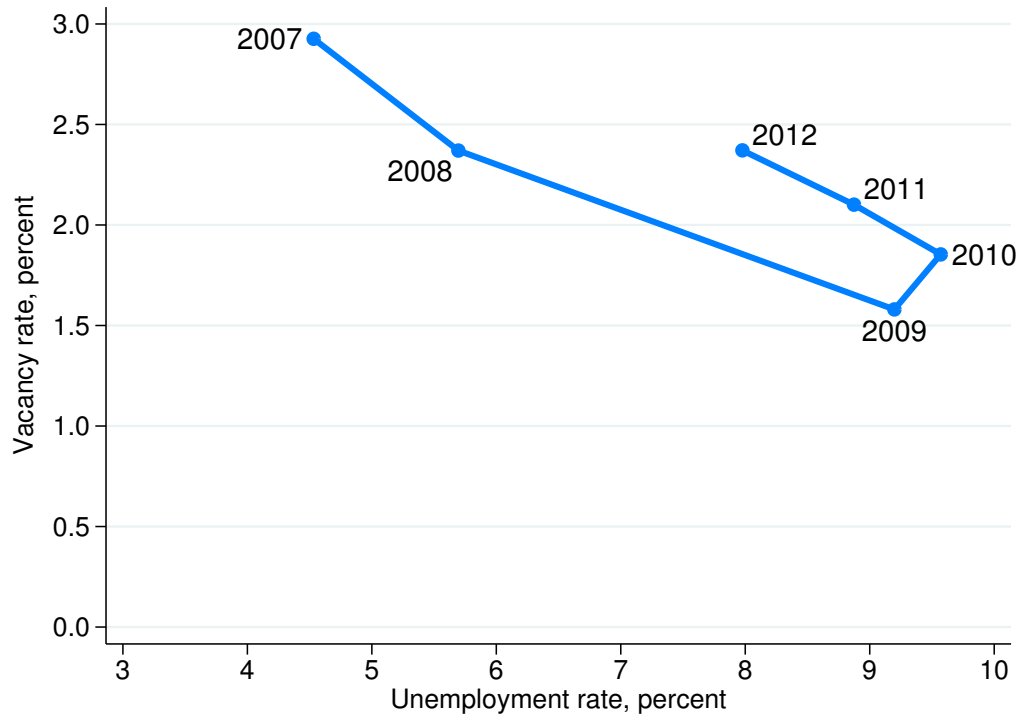


Figure 8: Beveridge Curve with the Standard Unemployment Rate Taken as the Measure of Jobseekers

Source: JOLTS and authors' calculations from Current Population Survey microdata. Annual averages of monthly data.

that includes job-to-job movers, one for short spans that excludes job-to-job movers, and one for long spans that excludes job-to-job movers. Movements in all three of these indexes in the crisis and its aftermath are smaller than the movements in unemployment in Figure 8, because unemployment rose much more than the other categories, and because unemployment tended to rise most in labor market statuses that have relatively low levels of matching efficiency—the crisis sent people into unemployment who were quite a bit harder to re-employ than those who quit jobs, to give a leading example. Including job-to-job movers further reduces movements of the job-seeking index, because the share of the population that is employed fluctuates less, in percentage terms, than the share that is unemployed. But the outward shift of the Beveridge curve in Figure 9 is still quite pronounced regardless of which index we use. Matching efficiency did decline substantially after 2007.

Figure 10 shows that almost all of the adverse shift of the Beveridge curve arose from trends that were in place prior to the crisis. It calculates the index of effective jobseeking using the trend coefficients from the model, which are mostly negative. With the adjustment, almost all of the outward shift disappears. In 2012, the labor market would have been on the




Figure 9: Beveridge Curve with the Index of Jobseekers Based on Detrended Matching Efficiency

Source: JOLTS and authors' calculations from Current Population Survey microdata. Annual averages of monthly data.

same Beveridge curve as in 2007 and 2008, had the downward trend in matching efficiency present from 2001 through 2007 not continued.

7 Conclusion

Many authors have demonstrated a decline in labor-market matching efficiency during the Great Recession and ensuing slump. With the exception of Veracierto's pioneering work, research has made the assumption that the measure of job-seeking volume is the stock of unemployed workers. But the Current Population Survey shows that only about a quarter of newly filled jobs involves hires of the unemployed. The remaining three-quarters come from out of the labor market or from job-to-job transitions. We develop a consistent approach to aggregation over heterogeneous categories of job-seekers, with a separate measure of matching efficiency for each category and a related measure of aggregate matching efficiency.

A second novel element in our work is to study the effectiveness of job search over spans greater than a month. Longer spans have two advantages: First, they lower the bias from misclassification, which tends to overstate job-finding rates measured as monthly transition




Figure 10: Beveridge Curve with the Index of Jobseekers Based on Matching Efficiency, Including Trends

Source: JOLTS and authors' calculations from Current Population Survey microdata. Annual averages of monthly data.

rates from jobseeking to employment. Second, they give less weight to transitory interim jobs, which appear to be an important part of the jobseeking process.

Our concept of matching efficiency combines the propensity of the members of a category of potential job-seekers to engage in active search with the per-period effectiveness of those active searchers. Absent direct measures of search effort, as in Krueger and Mueller (2011), we cannot break the two factors apart.

We confirm that matching efficiency has declined in some categories of unemployment, including permanent job loss, a category that rose substantially as a fraction of total unemployment in the Great Recession. Most of the decline is the continuation of a trend that has existed since 2001 and possibly earlier. Because such a large fraction of hiring occurs out of pools of job-seekers other than the unemployed, one important implication is that the decline in matching efficiency among the unemployed drove up the unemployment rate, but the labor market still generated large volumes of job-finding among groups not counted as unemployed.

References

- Abowd, John M. and Arnold Zellner, “Estimating Gross Labor-Force Flows,” *Journal of Business & Economic Statistics*, 1985, 3 (3), 254–283.
- Bachmann, Ronald and Mathias Sinning, “Decomposing the Ins and Outs of Cyclical Unemployment,” Technical Report DP No. 6362, IZA February 2012.
- Barlevy, Gadi, “Evaluating the Role of Labor Market Mismatch in Rising Unemployment,” *Economic Perspectives*, 3rd Quarter 2011, pp. 82–96. Federal Reserve Bank of Chicago.
- Barnichon, Regis and Andrew Figura, “The Determinants of the Cycles and Trends in U.S. Unemployment,” June 2012. Federal Reserve Board.
- Blanchard, Olivier J. and Peter A. Diamond, “The Cyclical Behavior of the Gross Flows of U.S. Workers,” *Brookings Papers on Economic Activity*, 1990, (2), 85–143.
- Daly, Mary C., Bart Hobijn, and Robert G. Valletta, “The Recent Evolution of the Natural Rate of Unemployment.,” 2011. IZA Discussion Paper No. 5832.
- , —, Ayşegül Şahin, and Robert G. Valletta, “A Search and Matching Approach to Labor Markets: Did the Natural Rate of Unemployment Rise,” *Journal of Economic Perspectives*, Summer 2012, 26 (3), 3–26.
- Davis, Steven J. and John Haltiwanger, “Labor Market Fluidity and Economic Performance,” September 2014. Forthcoming in the proceedings of the Jackson Hole Symposium, Federal Reserve Bank of Kansas.
- , R. Jason Faberman, and John C. Haltiwanger, “The Establishment-Level Behavior of Vacancies and Hiring,” *Quarterly Journal of Economics*, 2013, 128 (2), 581–622.
- , —, —, and Ian Rucker, “Adjusted Estimates of Worker Flows and Job Openings in JOLTS,” in Katharine G. Abraham, James R. Spletzer, and Michael J. Harper, eds., *Labor in the New Economy*, Chicago: University of Chicago Press, 2010, chapter 5.
- Elsby, Michael W.L., Bart Hobijn, and Ayşegül Şahin, “On the Importance of the Participation Margin for Market Fluctuations,” Working Paper 2013-05, Federal Reserve Bank of San Francisco February 2013.

- ____, _____, _____, and Robert G. Valletta, “The Labor Market in the Great Recession: An Update to September 2011,” *Brookings Papers on Economic Activity*, 2011, pp. 353–384.
- Estevão, Marcello and Christopher Smith, “Skill Mismatches and Unemployment in the United States,” January 2013. International Monetary Fund.
- ____ and Evidiki Tsounta, “Has the Great Recession Raised U.S. Structural Unemployment?,” May 2011. International Monetary Fund.
- Farber, Henry S. and Robert G. Valletta, “Do Extended Unemployment Benefits Lengthen Unemployment Spells? Evidence from Recent Cycles in the U.S. Labor Market,” Working Paper 19048, National Bureau of Economic Research, May 2013.
- Feng, Shuaizhang and Yingyao Hu, “Misclassification Errors and the Underestimation of the US Unemployment Rate,” *American Economic Review*, 2013, *103* (2), 1054–70.
- Flinn, Christopher J. and James J. Heckman, “Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?,” *Journal of Labor Economics*, 1983, *1* (1), 28–42.
- Fujita, Shigeru, “Effects of Extended Unemployment Insurance Benefits: Evidence from the Monthly CPS,” January 2011. Federal Reserve Bank of Philadelphia.
- ____ and Giuseppe Moscarini, “Recall and Unemployment,” June 2013.
- Hagedorn, Marcus, Fatih Karahan, Iourii Manovskii, and Kurt Mitman, “Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects,” October 2013. National Bureau of Economic Research Working Paper 19499.
- Hall, Robert E., “Lost Jobs,” *Brookings Papers on Economic Activity*, 1995, *1995* (1), 221–273.
- ____, “High Discounts and High Unemployment,” January 2014. Hoover Institution, Stanford University.
- ____, “Quantifying the Lasting Harm to the U.S. Economy from the Financial Crisis,” *NBER*, 2014. Forthcoming.

- Herz, Benedikt and Thijs van Rens, “Structural Unemployment,” July 2011. Universitat Pompeu Fabra.
- Hyatt, Henry R. and James R. Spletzer, “The Recent Decline in Employment Dynamics,” March 2013. Center for Economic Studies, US Census Bureau.
- Kolenikov, Stanislav, “Resampling variance estimation for complex survey data,” *Stata Journal*, 2010, 10 (2), 165–199.
- Krueger, Alan B. and Andreas I. Mueller, “Job Search and Job Finding in a Period of Mass Unemployment: Evidence from High-Frequency Longitudinal Data,” January 2011. Princeton University.
- , Judd Cramer, and David Cho, “Are the Long-Term Unemployed on the Margins of the Labor Market?,” *Brookings Papers on Economic Activity*, Spring 2014.
- Madrian, Brigitte C. and Lars John Lefgren, “An approach to longitudinally matching Current Population Survey (CPS) respondents,” *Journal of Economic and Social Measurement*, January 2000, 26 (1), 31–62.
- Nakajima, Makoto, “A Quantitative Analysis of Unemployment Benefit Extensions,” *Journal of Monetary Economics*, 2012, 59 (7), 686 – 702.
- Nekarda, Christopher J., “A Longitudinal Analysis of the Current Population Survey: Assessing the Cyclical Bias of Geographic Mobility,” May 2009. Board of Governors of the Federal Reserve System.
- Petrongolo, Barbara and Christopher A. Pissarides, “Looking into the Black Box: A Survey of the Matching Function,” *Journal of Economic Literature*, 2001, 39 (2), 390–431.
- Poterba, James M. and Lawrence H. Summers, “Reporting Errors and Labor Market Dynamics,” *Econometrica*, 1986, 54 (6), pp. 1319–1338.
- Rao, J.N.K., C.F.J. Wu, and K. Yue, “Some recent work on resampling methods for complex surveys,” *Survey Methodology*, 1992, 18, 209–217.
- Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante, “Mismatch Unemployment,” Working Paper 18265, National Bureau of Economic Research August 2012.

Valletta, Rob and Katherine Kuang, “Is Structural Unemployment on the Rise?” *Federal Reserve Bank of San Francisco Economic Letter*, November 2010, 2010 (34), 1–5.

Veracierto, Marcelo, “Worker flows and matching efficiency,” *Economic Perspectives*, 4th quarter 2011, pp. 147–169.

A Related Research

Veracierto (2011) introduced the basic idea of including people other than the unemployed in the calculation of matching efficiency. He makes a compelling case that the movements of aggregate unemployment cannot be understood in the DMP framework—especially with respect to the matching function—without considering the role of individuals who are classified as out of the labor market. These people are neither working nor engaging in the specific job-seeking activities in the four weeks prior to the CPS interview that would place them in the category of unemployment. The striking fact is that, after correcting in the standard way for erroneous transitions, the CPS reveals that the number of people classified as out of the labor force in one month who are employed in the next month is always greater than the number moving from unemployment to employment. In normal times, using the obvious notation, the NE flow is almost double the UE flow.

Flinn and Heckman (1983) observe that the natural definition of unemployment is that a non-working individual's transition hazard into employment exceeds a threshold value. By that criterion, it seems likely that a non-trivial fraction of those the CPS classifies as out of the labor force (N) are actually unemployed. But the overall NE hazard in normal times is far lower than the UE hazard—5 percent per month compared to 27 percent, so it is clear that the U category in general satisfies the Flinn-Heckman criterion.

The BLS publishes data on broader definitions of unemployment. It is an interesting question but outside the scope of this paper whether a systematic application of the Flinn-Heckman principle might result in a definition of unemployment that captured the great majority of non-workers with high job-finding hazards while excluding those with low hazards. Such a definition would fit the matching function framework nicely.

Veracierto (2011) proposes a simple way around this issue that incorporates those classified as out of the labor force without identifying the individuals with high NE hazards. A brief discussion in Petrongolo and Pissarides (2001), p. 403, anticipates Veracierto's approach. He uses the ratio of the NE hazard to the UE hazard to weight those classified in N. The resulting figure is interpreted as the effective number of job-seekers in the N category. The total number of job-seekers is the number in U plus the weighted number in N. This figure—interpreted as comprehensive unemployment—is the input to the matching function in a DMP model that takes account of the high incidence of job-seeking in the N

category. Veracierto finds (see his figure 36) that matching efficiency was flat before the Great Recession, then declined about 15 percent during the recession.

Our analysis differs from Veracierto's both in the definition of matching efficiency and in the level of disaggregation. Veracierto assumes that unemployed workers and nonparticipants have equal matching efficiency conditional on a given level of search intensity but that nonparticipants have lower search intensity. By contrast, we do not distinguish between matching efficiency and search intensity for a given type of worker and instead estimate an efficiency parameter for each type that combines matching efficiency and search intensity. In addition, our analysis includes job-to-job transitions and further disaggregates workers by their reason for unemployment and by observable characteristics. Our model thus provides a unified treatment of the calculation of aggregate matching efficiency when all people in the economy of working age are potentially job seekers.

Barnichon and Figura (2012) also estimate matching efficiency while allowing heterogeneity across workers in demographics, distinguishing between reasons for unemployment, and including nonparticipants in the analysis. However, they assume that the matching function applies only to unemployed workers and do not consider job-to-job transitions.

Fujita and Moscarini (2013) study the effect of recalls by unemployed workers' former employers on transition rates and the matching function. They show that if the matching function describes only matches between jobseekers and new employers—not recalls—then matching efficiency is estimated to have declined much more during the Great Recession. Key to their result is that workers on temporary layoffs are not the only ones who experience recalls; about 20 percent of workers who report that they permanently lost their jobs are nonetheless eventually recalled. In our work, we disaggregate workers by their reason for unemployment but do not attempt to distinguish between matches with new employers and recall by the previous employer. Thus, in our specification, a group that is more likely to be recalled will have a higher matching efficiency.

Barlevy (2011) calculates the decline in matching efficiency from the shift in the Beveridge curve, on the assumptions that the separation rate remains unchanged and that unemployment is at its stochastic equilibrium. This analysis depends only on the unemployment rate, not on the number of nonparticipants, job-to-job transitions, or changes in the composition of the unemployed.

Bachmann and Sinning (2012) measure the effects of compositional changes on labor force transition rates without relating these findings to matching efficiency. They find that changes in composition reduce the cyclical nature of inflows to unemployment and raise outflows from unemployment early in recessions but reduce outflows later in recessions.

Some papers discuss the decline in matching efficiency, or, equivalently, the outward shift of the Beveridge curve, as the result of a variety of forces. Some, such as Daly, Hobijn, Şahin and Valletta (2012), frame the subject within the more general issue of a possible increase in the natural rate of unemployment. Only part of their discussion relates to changes in matching efficiency. The paper identifies two factors that may have reduced match efficiency since the Great Recession: mismatch and more generous unemployment benefits.

Şahin, Song, Topa and Violante (2012) find that mismatch across industries and occupations accounts for at most one-third of the increase in unemployment during the Great Recession, while geographic mismatch is insignificant. Herz and van Rens (2011) likewise find modest effects of mismatch across industries and very small effects of mismatch across states, while Estevão and Tsounta (2011) find substantial skill mismatches but argue that changes in migration rates and dispersion in unemployment across states are evidence of geographic mismatch as well. These studies all measure mismatches by the distribution of unemployed workers and jobs across distinct markets defined by locations, industries, or occupations. Estevão and Smith (2013) measure skill mismatches in a different way, by imputing wages for labor force participants based on their observed characteristics; if mismatch is low and unemployment is mainly due to low quality of unemployed workers, unemployed workers will have relatively low imputed wages, while if mismatch is high, unemployed workers will have relatively high imputed wages. Consistent with the papers that look at mismatch across distinct markets, Estevão and Smith (2013) find evidence of an increase in mismatch during the recession.

A number of papers, including Daly, Hobijn and Valletta (2011), Fujita (2011), Nakajima (2012), and Valletta and Kuang (2010), culminating in Farber and Valletta (2013), find that extended unemployment benefits raised the unemployment rate by an amount ranging from a few tenths of a percentage point to one point. However, Hagedorn, Karahan, Manovskii and Mitman (2013) argue that many of these analyses do not account for the effect of unemployment benefits on firms' incentive to create jobs and that a research design that

accounts for such effects finds a much larger impact from unemployment benefits. Hall (2014b) discusses their paper at greater length.

Davis, Faberman and Haltiwanger (2013) provide convincing evidence that the matching function involves inputs apart from the stocks of unemployment and vacancies. In the micro data from JOLTS, they show that the job-filling rate for vacancies is dramatically higher in firms that are growing than in firms with constant employment, a contradiction to the hypothesis that only unemployment and vacancies determine hiring rates. They lack any direct measures of the other inputs, but construct an indirect measure from the JOLTS data that eliminates most of the apparent decline in matching efficiency. They do not consider the topic of this paper, the importance of job-seekers who are not counted as unemployed. Their results fit nicely with ours, in the sense that one reasonable interpretation of the variations in matching efficiency that we measure is exactly the combined effect of the omitted inputs to the matching process that they consider.