

NBER WORKING PAPER SERIES

MEASURING JOB-FINDING RATES AND MATCHING EFFICIENCY WITH HETEROGENEOUS
JOBSEEKERS

Robert E. Hall
Sam Schulhofer-Wohl

Working Paper 20939
<http://www.nber.org/papers/w20939>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2015, Revised July 2017

The Hoover Institution supported Hall's research. The research is also part of the NBER's Economic Fluctuations and Growth Program. We thank Suyoun Han for excellent research assistance and Christopher Nekarda for sharing his Stata code for longitudinal matching of observations in the Current Population Survey. We have benefited from comments from the editor, referees, and numerous conference and seminar participants, especially Steve Davis, who also provided us with the measures of vacancy duration and recruiting intensity from his work with Jason Faberman and John Haltiwanger, updated. A previous version of this paper was written while Schulhofer-Wohl was employed at the Federal Reserve Bank of Minneapolis. The views expressed here are those of the authors and not necessarily those of the National Bureau of Economic Research, the Federal Reserve Bank of Chicago, the Federal Reserve Bank of Minneapolis, or the Federal Reserve System. Emails: re- hall@stanford.edu and sam@frbchi.org. Data and program files for replicating our work are available at Hall's website, stanford.edu/~rehall

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Robert E. Hall and Sam Schulhofer-Wohl. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Jobseekers
Robert E. Hall and Sam Schulhofer-Wohl
NBER Working Paper No. 20939
February 2015, Revised July 2017
JEL No. E24,J63

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 an 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 smoothly over the period from 2001 through 2013. Measures of matching efficiency that neglect heterogeneity among the unemployed and also neglect jobseekers other than the unemployed suggest a 23 percent decline in efficiency between 2007 and 2009. We demonstrate that essentially all of this apparent decline results from changes in the composition of jobseekers rather than any true movement in efficiency. We also develop a new approach to measuring matching rates that avoids counting short-duration jobs as job-seeking successes.

Robert E. Hall
Hoover Institution
Stanford University
Stanford, CA 94305-6010
and NBER
rehall@gmail.com

Sam Schulhofer-Wohl
Economic Research Department
Federal Reserve Bank of Chicago
230 South LaSalle Street
Chicago, IL 60604-1413
sam@frbchi.org

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. High unemployment from late 2008 until 2013 generated concern that the U.S. economy’s normal unemployment rate rose from the turmoil of the financial crisis. We show that disaggregated measures of matching efficiency did not have large declines after the crisis. Rather, the composition of unemployment shifted dramatically toward groups with chronically lower matching efficiency.

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. In our framework, matching efficiency is a set of multiplicative shifters of the job-seeking population, analogous to factor-augmenting productivity indexes in production theory. We measure matching efficiency using standard ideas from production theory. In the DMP model, a decline in matching efficiency changes labor-market equilibrium in the direction of higher unemployment. Proper measurement of matching efficiency is a crucial starting point for understanding the sources of episodes of high unemployment.

This paper measures matching efficiency. It does not attempt to explain why matching efficiency changes over time, in response to its economic determinants.

Section 1 of the paper lays out our theory of matching with heterogeneous jobseekers. We consider the case where a single index of labor-market tightness governs the jobfinding rates of all categories of jobseekers, but the elasticity of the jobfinding rate with respect to tightness varies across categories. Each category also has a distinct parameter measuring matching efficiency. We show that this assumption about jobfinding rates implies a matching function with standard properties, including constant returns to scale. In the case of a single category of jobseeker, the matching function is Cobb-Douglas.

Section 2 considers the application of our theory to data from the Current Population Survey public use sample. We consider 16 categories of jobseekers: two for people out of the labor force, one for workers, and 13 for the unemployed based on the duration of unemployment and the reasons for unemployment reported in the CPS, such as “lost permanent job.” Because people who are out of the labor force or looking for work may take jobs that are extremely brief, measuring the jobfinding rate as the probability of taking a new job from one month to the next overstates the success of jobseekers in finding durable jobs. We examine the probability of employment not only one month later but also two, three, and up to 15 months later. We characterize these probabilities as logit functions of personal characteristics and time effects, separately by category of jobseeker and the number of months the probability spans. Overall, our specification captures heterogeneity by (1) standard de-

mographic characteristics, (2) desire and availability for work among people out of the labor force, (3) reason for unemployment and duration, for those unemployed, and (4) the number of months over which the probabilities are measured.

Section 3 describes the estimated logit probabilities along a number of dimensions. The logits are used to construct jobfinding rates that hold demographics constant within the 16 job-seeking categories described above. We find a downward trend of jobfinding probabilities when measured over one to three months, but almost no trend when measured over 12 to 15 months. We find that the probability that a jobseeker will be employed 12 to 15 months later is hardly higher than the probability she will be employed one to three months later, which we interpret as demonstrating the key importance of quite short jobs in the experiences of jobseekers.

Section 4 connects independent data from the Job Openings and Labor Turnover Survey on labor market tightness, measured as the average duration of a vacancy, with our measures of demographically adjusted jobfinding rates from the logit models. Our matching function implies that the elasticity of the job-finding rate with respect to tightness is a constant over time within a category but varies across categories. We estimate these elasticities with an instrumental variable.

In Section 5, we adjust the job-finding rates by the estimated elasticities applied to the tightness measure, to form indexes of category-specific indexes of matching efficiency. We find a mild downward trend at close to a constant rate in most categories, with no special decline during the contraction from 2007 to 2010. We construct an aggregate index of matching efficiency by applying fixed weights across categories. It shares the downward trend of the category-specific indexes and, like them, has no special decline between 2007 and 2010. We compare our index with one constructed by similar principles, except that only one category of jobseeker is included, defined as unemployment in all categories. That index declines by almost 50 percent from 2007 to 2010. Essentially all of the decline arises from neglect of heterogeneity.

We conclude that it is essential to recognize heterogeneity among jobseekers and include all jobseekers, not just the unemployed, in measuring and understanding matching efficiency.

One of the key points of this paper is that the majority of jobseekers are not counted as unemployed, but rather as out of the labor force or employed. Despite Blanchard and Diamond's (1990) emphasis on this point, most analysis of the U.S. labor market in the matching-function framework has taken unemployment as the measure of job-seeking in the population. An important exception is Veracierto (2011), a paper that we build on.

Barnichon and Figura (2015b) preceded us in measuring matching efficiency with heterogeneous unemployed jobseekers. Our main contributions building on their paper are: (1) We consider the remaining groups responsible for the great majority of jobfinding success—we distinguish two categories of people recorded as out of the labor force, one with low jobfinding

propensities and another with higher propensities, following Barnichon and Figura (2015a). Most importantly, we include currently employed jobseekers, who are hoping for a job-to-job transition. (2) We focus on measuring job-seeking success over longer spans of time, up to the maximum possible in the CPS of 15 months. (3) We measure labor-market tightness as the ratio of vacancies to hires, rather than as the ratio of vacancies to unemployment. Thus we avoid using data on the right-hand side of our basic estimating equation that also appear on the left-hand side. (4) We introduce a new class of matching functions, suited to dealing with multiple categories of jobseekers.

1 Matching Functions with Heterogeneous Job-seekers

For the purposes of this paper, a matching function is a function $m(P, V)$, increasing and weakly concave in a vector of types of jobseekers P and the scalar number of vacancies V . $H = m(P, V)$ is the flow of new hires emerging from the matching process. H is analogous to output and P and V are analogous to factor inputs, so m is analogous to a production function. Let $T = V/H$, the ratio of vacancies to hires, which is the average duration of a vacancy. T measures the tightness of the labor market.

Assumption. Common pools of vacancies and competing jobseekers:

All types of jobseekers have success rates that depend on the same scalar measure of tightness, T .

Job-seekers of type i have an increasing job-seeking success hazard $\phi_i(T)$. The flow of new hires is

$$H = \sum_i \phi_i(V/H) P_i. \quad (1)$$

The unique solution to this equation defines the matching function $H = m(P, V)$. The assumption that the success hazard depends on T rather than separately on P and V implies that the matching function has constant returns to scale.

In the standard setup, the functions $\phi_i(T)$ are the same for all types. In that case, the matching function solves

$$H = \phi(V/H) \sum_i P_i, \quad (2)$$

which implies a standard matching function $H = m(V, \sum_i P_i)$.

We assume that the elasticity within each category is constant:

$$\phi_i(T) = \gamma_i T^{\eta_i} \quad (3)$$

The parameter γ_i conveys the level of matching efficiency for type i , because it multiplies the rate at which job-seekers of type i are converted into hires. The elasticity η_i conveys the

response of the matching rate to market tightness. One source of differences in the elasticity would be differences in the responsiveness of search effort to tightness. Our empirical work finds unambiguous differences in the elasticities.

An index of weighted average matching efficiency at a reference level of tightness \bar{T} and a reference set of population shares s_i is

$$\mu = \sum_i s_i \gamma_i \bar{T}^{\eta_i}. \quad (4)$$

For comparison with other estimates of labor-market matching functions, we note that the elasticity of the matching flow rate with respect to V is

$$\frac{\sum_i \gamma_i \eta_i T^{\eta_i} s_i}{\sum_i \gamma_i (1 + \eta_i) T^{\eta_i} s_i}. \quad (5)$$

In the standard case of only one kind of jobseeker, the elasticity is η , a constant across i , and the matching function is Cobb-Douglas with elasticity $\frac{\eta}{1+\eta}$.

Appendix A provides a detailed comparison of the approach we use, based on measuring tightness as the duration of a vacancy, with the standard DMP approach, based on measuring tightness as the vacancy/unemployment ratio. The two approaches are fully isomorphic with a single type of jobseeker. We demonstrate that our setup delivers jobfinding rates that are driven by a single index, and that the index is the duration of a vacancy, which is the same for all types.

2 Measurement of Job-Finding Rates in the CPS Data

We use data from the monthly CPS for November 1999 through March 2015. These data permit the calculation of jobfinding rates for individuals who started their searches in the years 2001 through 2013.

2.1 Labor-market status

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 also divide those out of the labor force into two categories. One is those who answer “no” to the question, “Do you want a job now, either full or part-time?” or who answer “yes” but then indicate they are not currently available. The other category is those who want a job and are available. Barnichon and Figura (2015a) found large differences in jobfinding rates of people classified as out of the labor force between those wanting work and those not wanting work.

We derive a total of 16 labor-market statuses. The three basic statuses are (1) out of labor force, (2) unemployed, and (3) employed. We consider two statuses within the basic status *out of labor force*:

- *Do not want work*: people who did not satisfy the CPS definition of either employed or unemployed and who did not want work or were not available to work
- *Want work but not looking*: people who did not satisfy the CPS definition of either employed or unemployed and who wanted work and were available to work

For those unemployed for three weeks or less, we define six categories:

- *Recently furloughed* : 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 furlough, 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 labor force*: unemployed people who have never worked and who started looking for work within the past three weeks.
- *Recently re-entered labor force*: 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 unemployment statuses parallel those above, with duration of unemployment to date of 4 to 26 weeks:

- *On furlough for months*
- *Lost permanent job months ago*

- *Temp job ended months ago*
- *Quit months ago*
- *Entered labor force months ago*
- *Re-entered labor force months ago*

The last unemployment 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.

Finally, we use a single category for employed:

- *Employed*: employed people

Ideally, we would distinguish between employed people by duration of employment, because people who have been in their jobs only a short time will be more likely to make job-to-job moves if jobs are heterogeneous in duration or if workers are heterogeneous in their propensity to move. However, the CPS data do not contain sufficient information on job tenure to allow us to draw such distinctions.

2.2 Time span for measuring jobfinding success

The standard concept of a jobfinding rate is the probability that a jobseeker 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.

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 in March 2009, on account of the earlier loss of a permanent job, and unemployed as well in April 2010. The span between these observations in our sense would then be 13 months.

It is important to understand that span is different from other measures, such as the duration of unemployment. The beginning of a span is not necessarily in the month the person entered the CPS or the month when a spell of unemployment began. In the 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. And the

<i>Calendar month</i>	<i>CPS month</i>	<i>Span, months</i>	<i>Unemployment duration, months</i>
November 2008			0
December 2008			1
January 2009			2
February 2009	1		3
March 2009	2	0	4
April 2009	3	1	5
May 2009	4	2	6
June 2009		3	7
July 2009		4	8
August 2009		5	9
September 2009		6	10
October 2009		7	11
November 2009		8	12
December 2009		9	13
January 2010		10	14
February 2010	13	11	15
March 2010	14	12	16
April 2010	15	13	17
May 2010	16		18
June 2010			19

Table 1: Example of CPS Survey Months, an Unemployment Spell Beginning in November 2008, and a Span Beginning in March 2009

person might 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.

A spell of unemployment may well be contained within a span. We observe people unemployed when they enter the CPS, employed briefly, then unemployed, and then employed late in the span. Turnover within spans has a central role in our empirical analysis.

Over the spans, we focus on the experiences of people who began the span in a given labor-market status, such as *recently quit*. We measure the probability that such a person would be employed, say, 12 months later.

Longer spans matter for measuring jobfinding success because jobseekers may find brief jobs, lasting only a few weeks or a month or two—see Hyatt and Spletzer (2013). 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. Job-seekers 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.

Because the CPS interviews households for 4 consecutive months, skips the next 8 months, then interviews again for 4 months, each person covered for every scheduled interview contributes 6 observations spanning single months, 4 spanning 2 months, 4 spanning 12 months, and one spanning 15 months, to give a few examples. In principle, we can study job-seeking spans of 1, 2, 3, 9, 10, 11, 12, 13, 14, and 15 months. For simplicity, we omit the 9-, 10- and 11-month spans and focus on the short spans from 1 through 3 months and the long spans from 12 through 15 months.

Our choice to examine short and long spans is aimed at extracting useful information about job-finding success from the limited data in the CPS. The ideal data would contain accurate month-by-month records of labor-market status over many years. Such data would support a model that showed how individuals who start from a given status find jobs, lose or leave them, search more, take subsequent better jobs, and ultimately settle into durable job matches or non-work activities. The CPS reports in only eight separate months over a period of a bit more than a year. Our short-span measure uses just the initial four months, whereas the long-span measure considers what has happened about a year after the conditioning month. Both measures describe experience that is influenced by both jobfinding success and separations. See Krueger, Cramer and Cho (2014) for further discussion of the benefits of including longer spans in the study of jobfinding success.

2.3 Specification of the jobfinding function

In the empirical implementation of our setup, the jobfinding 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

$$f_{i,t,\tau,x} = \frac{\exp(\kappa_{i,t,\tau} + x'\beta_{i,\tau})}{1 + \exp(\kappa_{i,t,\tau} + x'\beta_{i,\tau})}, \quad (6)$$

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.

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 jobfinding 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 study jobfinding success conditional on standard observable demographic characteristics and on the initial status. We find large differences in jobfinding rates by initial status. Our results would have the cleanest interpretation if the demographic characteristics accounted for all the heterogeneity in each group defined by initial status. We do not make that claim—people laid off during a recession, for example, differ in some respects not captured by their observable demographic characteristics from people laid off during a boom. We find that overall labor-market tightness accounts for most of the large movements in jobfinding rates around trend within each initial-status group. This finding supports the view that our breakdown by 16 groups captures most of the heterogeneity among jobseekers.

2.4 Processing the CPS data

We match respondents across months using the method of Nekarda (2009). His 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 jobfinding 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 jobfinding rates across long time spans because errors in recording race, sex, and age during intervening months are less likely to break the match. Table 7 in Appendix B shows the success rates for the matching process.

We remove high-frequency, likely spurious transitions between unemployment and non-participation following Elsby, Hobijn and Şahin (2015b). Specifically, if a respondent is out

of the labor force, unemployed, and out of the labor force again in three consecutive months, we recode the middle month to *want work*, if the respondent wanted to work in either the first or third month; if not, we recode 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 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.

A substantial literature describes reporting errors in the CPS and similar longitudinal surveys. Even after following the Elsby et al. (2015b) procedure, 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.

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 as an indicator of the time that has elapsed since the individual has held a job more durable than an interim job.

We use a block bootstrap procedure to compute standard errors that account for the CPS sampling design and for our use of multiple spans for the same individual, both of which induce cross-observation correlations that would be difficult to account for analytically. 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 100 bootstrap samples by resampling households

with replacement within each state-month, thus reproducing the stratification of the CPS by state. Each resampling follows all household members through all subsequent appearances in the CPS, so we reproduce the correlations across members and over time within each household that result from the dataset’s rotating panel structure and from our use of multiple spans per individual. 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.

In both the original sample and each bootstrap sample, we reweight the data in a standard way to account for attrition from the survey, though we found that this adjustment had little effect on the estimated jobfinding rates. Appendix B provides details on the adjustment.

3 Estimated Job-Finding Rates

3.1 Job-finding rates implied by the logit specification

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 jobfinding rates standardized for the changing composition of the labor force. We standardize 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 2013, we calculate the fitted jobfinding 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. Appendix F describes how we handle technical issues that arise for a few status-month-span cells that contain few or no observations.

3.2 Employment probabilities by initial status

Table 2 summarizes our findings for employment probabilities conditional on originating in each of the job-seeking statuses. The left panel shows the probabilities averaged over short spans and the right panel over long spans. The table compares the employment probabilities in 2013 with those in 2003 because the business cycle was in a similar phase in the two years.

The employment probabilities in Table 2 vary over a wide range across the initial statuses. Thus our decomposition by initial status captures a great deal of heterogeneity in jobfinding

<i>Initial status</i>	<i>Average employment probability, months 1 to 3</i>			<i>Average employment probability, months 12 to 15</i>		
	2003	2013	<i>Ratio</i>	2003	2013	<i>Ratio</i>
Do not want work (Standard error)	5.7 (0.1)	4.5 (0.0)	0.78 (0.01)	11.8 (0.2)	9.9 (0.2)	0.84 (0.02)
Want work but not looking (Standard error)	16.9 (0.4)	14.9 (0.3)	0.88 (0.03)	32.3 (0.8)	30.8 (0.7)	0.95 (0.03)
Employed (Standard error)	5.2 (0.1)	4.5 (0.0)	0.87 (0.01)			
Recently furloughed (Standard error)	59.8 (1.3)	59.2 (1.4)	0.99 (0.03)	64.7 (2.0)	68.7 (1.7)	1.06 (0.04)
Recently lost permanent job (Standard error)	34.6 (1.4)	35.3 (2.0)	1.02 (0.07)	67.9 (2.2)	63.5 (2.4)	0.94 (0.04)
Temp job recently ended (Standard error)	44.2 (2.4)	40.3 (2.4)	0.91 (0.07)	62.5 (3.5)	60.5 (3.4)	0.97 (0.08)
Recently quit (Standard error)	42.9 (2.2)	42.6 (2.3)	0.99 (0.08)	64.5 (3.6)	65.9 (3.7)	1.02 (0.08)
Recently entered labor force (Standard error)	30.1 (2.7)	20.8 (1.8)	0.69 (0.09)	51.0 (4.4)	39.5 (3.6)	0.77 (0.09)
Recently re-entered labor force (Standard error)	35.0 (1.3)	31.3 (1.3)	0.89 (0.05)	50.4 (2.3)	48.7 (2.1)	0.97 (0.06)
On furlough for months (Standard error)	46.6 (1.5)	48.9 (1.5)	1.05 (0.05)	57.9 (2.3)	60.2 (2.4)	1.04 (0.06)
Lost permanent job months ago (Standard error)	26.0 (0.8)	26.7 (1.0)	1.03 (0.05)	62.7 (1.4)	57.8 (1.6)	0.92 (0.03)
Temp job ended months ago (Standard error)	30.2 (1.5)	28.9 (1.5)	0.96 (0.07)	54.3 (2.7)	54.3 (2.5)	1.00 (0.07)
Quit months ago (Standard error)	34.8 (1.4)	31.5 (1.6)	0.91 (0.06)	58.7 (2.7)	57.2 (3.0)	0.97 (0.06)
Entered labor force months ago (Standard error)	21.6 (1.7)	15.6 (1.0)	0.72 (0.07)	44.3 (3.1)	44.6 (2.7)	1.01 (0.09)
Re-entered labor force months ago (Standard error)	28.1 (0.9)	24.9 (0.9)	0.88 (0.04)	46.8 (1.6)	45.2 (1.6)	0.97 (0.05)
Long-term unemployed (Standard error)	19.8 (0.7)	16.4 (0.5)	0.83 (0.04)	43.2 (1.4)	40.4 (1.0)	0.93 (0.04)

Table 2: Standardized Subsequent Employment Probabilities for Short and Long Spans, 2003 and 2013, with Growth Ratio

rates. Not including the employed, for whom we look at the probability of changing jobs, the lowest jobfinding rate is for people starting in the status *do not want work*. In 2013, their short-span subsequent employment probability was 4.5 percent and their long-span rate was 9.9 percent. Most people classified as out of the labor force and not wanting work (or not being available for work) remain in non-market activities from one year to the next.

A striking feature of Table 2 is that the long-span jobfinding rates for workers on furlough (laid off with expectation of recall) are barely higher than the short-span jobfinding rates for these workers. This result implies either that laid-off workers have a very high separation rate upon reemployment or that, after a few months of unemployment, laid-off workers' reemployment hazard is very low. The latter implication is consistent with the findings of Katz (1986) and Fujita and Moscarini (2013) of strong negative duration dependence in recall probabilities for laid-off workers.

The long-term unemployed had short-span re-employment success rates of only 16 percent in 2013. Over the longer span of 12 to 15 months after the conditioning month (which is itself at least 6 months after the job loss), 40 percent of 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 2003, a year of somewhat lower overall unemployment than 2013. Our research deals with only the outflow rate from long-term unemployment. An understanding of the high levels of long-term unemployment following the crisis of 2008 would require a study of inflow rates to unemployment, a subject complementary to the subject of this paper—see Ahn and Hamilton (2016) and Hall and Schulhofer-Wohl (2017).

Entrants and re-entrants to the labor force 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 jobseekers. Losers of permanent jobs include those often termed *displaced*—that is, losers of high-tenure, seemingly secure jobs with high match surpluses arising from job-specific human capital.

3.3 Trends in jobfinding rates

Table 2 shows that there is generally a downward trend in jobfinding rates over the period we study. Only three of the initial status categories had short-span finding rates that were higher in 2013 compared to 2003. The trend in long-span finding rates is less frequently and less strongly downward. Only two of the initial status categories had 2013 rates below 90 percent of their 2003 rates. Earlier we noted that the long-span finding rate gives less weight to short jobs, so the finding of larger declines in short-span finding rates points in the direction of a role for a decline in the incidence of short jobs. This paper does not consider the incidence directly, because we do not study job separations. Davis and Haltiwanger




Figure 1: Estimated Job-Finding Probabilities for Losers of Permanent Jobs

(2014) called attention to the decline in labor turnover in recent years. Because short jobs account for the great majority of separations, a decline in overall separations necessarily involves a diminished incidence of short jobs.

3.4 The role of the span between initial observation and subsequent jobfinding success

Figure 1 shows the mix-adjusted estimated jobfinding probabilities for one important initial status, *recently lost permanent job*, separately by the span between initial observation and subsequent employment status. The data are annual averages of monthly calculations. The horizontal axis gives the date of search, corresponding to the calendar month of the CPS. Each curve is labeled on the right by the number of months in the future of the employment status of the respondent—for example, in 2005, the curve labeled 12 refers to survey responses recorded during 2005 referring to employment in 2006. The lowest curve is the probability that a person who lost a permanent job in the past three weeks 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 2013.

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 34 percent and the two-month proba-

bility was 43 percent. If the monthly jobfinding rate was truly 34 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 $0.34 + (1 - 0.34) \times 0.34 = 0.56$, far above the actual value of 0.43. The remaining curves in Figure 1 lie even closer to each other, so the property is even more acute for longer spans. A more refined benchmark recognizes that the average separation rate from jobs in general in the CPS is about 6 percent per month. With this factor taken into consideration, the probability of being employed in the second month would be 0.54, still well above the actual value.

Table 3 and Table 4 summarize the similar results for the 16 initial statuses, for the year 2007, before the dislocation of the financial crisis and deep recession. The subsequent employment 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 people in that status in a given month who were employed in the later months of the CPS interview schedule. For example, 4.1 percent of those in the *do not want work* status in a given month were employed in the following month and 11.9 percent 15 months later. The row labeled *Benchmark* is the projected percentage if the jobfinding rate for month 1 applies in all the later months, along with a monthly probability of 6 percent that any job found ends in a subsequent month and the worker cycles back to the status named at the left. 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 furloughed* status, which has a high one-month jobfinding rate of 56.0 percent, the benchmark would have 90.3 percent back at work 15 months later, but in fact, only 61.6 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.

The leading explanation for the shortfall of longer-span jobfinding rates from a benchmark based on average separation probabilities for jobs is that the separation rate is much higher for new jobs than for jobs in general. The separation rate needed to rationalize the low observed two-month success rate for job search while unemployed is 40 percent per month. Hyatt and Spletzer (2013) provide evidence from a variety of sources on the high incidence of short-duration jobs, including quarterly administrative data only recently available for this purpose. There is considerable remaining uncertainty about the monthly separation rate, but rates in the range of 40 percent per month are not ruled out. Another data source widely used to measure separation rates is the CPS tenure supplement, which collects data on the distribution of tenure among workers. The large number of workers with very low tenure implies high monthly separation rates.

Krueger et al. (2014) recently documented the shortfall in longer-span jobfinding rates for the unemployed. They studied long-term unemployment. They concluded, "...the long-term unemployed face difficulty regaining full-time, steady work over the longest period we can

<i>Initial status</i>		<i>Percent employed as of a later month</i>						
		<i>Months later</i>						
		1	2	3	12	13	14	15
Do not want work	Actual	4.1	5.6	6.5	10.9	11.2	11.6	11.9
	(Standard error)	(0.0)	(0.1)	(0.1)	(0.1)	(0.1)	(0.2)	(0.2)
	Benchmark	4.1	7.9	11.2	29.5	30.6	31.7	32.6
Want work but not looking	Actual	14.7	18.9	21.0	30.9	31.3	31.8	30.4
	(Standard error)	(0.3)	(0.4)	(0.6)	(0.7)	(0.7)	(0.9)	(1.1)
	Benchmark	14.7	26.3	35.6	66.6	67.5	68.2	68.8
Employed	Actual	56.0	64.9	64.9	62.2	60.2	58.2	61.6
	(Standard error)	(1.4)	(1.6)	(2.3)	(1.6)	(2.3)	(2.2)	(2.8)
	Benchmark	56.0	77.3	85.4	90.3	90.3	90.3	90.3
Recently furloughed	Actual	33.7	42.9	46.5	66.2	62.5	59.7	61.7
	(Standard error)	(1.7)	(2.1)	(2.7)	(2.2)	(2.5)	(3.0)	(3.8)
	Benchmark	33.7	54.0	66.2	84.7	84.8	84.8	84.8
Recently lost permanent job	Actual	42.1	54.1	49.1	59.9	61.2	66.2	56.9
	(Standard error)	(2.0)	(3.2)	(4.5)	(3.4)	(4.1)	(4.9)	(6.6)
	Benchmark	42.1	64.0	75.3	87.5	87.5	87.5	87.5
Temp job recently ended	Actual	40.3	51.7	58.1	69.1	64.1	67.5	58.8
	(Standard error)	(1.9)	(2.4)	(3.6)	(2.5)	(2.8)	(3.9)	(4.2)
	Benchmark	40.3	62.0	73.6	87.0	87.0	87.0	87.0
Recently quit	Actual	29.3	28.8	25.4	37.4	41.9	37.8	43.7
	(Standard error)	(2.5)	(3.2)	(3.1)	(4.0)	(4.1)	(5.3)	(7.6)
	Benchmark	29.3	48.2	60.5	82.5	82.7	82.8	82.9
Recently entered labor force	Actual	35.5	44.1	43.7	52.4	56.0	56.5	57.1
	(Standard error)	(1.3)	(1.7)	(2.3)	(2.3)	(2.3)	(3.1)	(3.6)
	Benchmark	35.5	56.2	68.4	85.4	85.5	85.5	85.5

Table 3: Subsequent Employment Probabilities by Initial Status, Actual and Benchmark, 2007: Out of Labor Force, Employed, and Recently Unemployed

benchmark1,2

		<i>Percent employed as of a later month</i>						
		<i>Months later</i>						
<i>Initial status</i>		1	2	3	12	13	14	15
Recently re-entered labor force	Actual	42.7	51.3	59.1	49.9	54.5	63.1	63.7
	(Standard error)	(1.6)	(2.1)	(3.2)	(2.5)	(3.0)	(3.5)	(4.2)
	Benchmark	42.7	64.6	75.9	87.7	87.7	87.7	87.7
On furlough for months	Actual	22.9	31.6	37.8	58.8	58.9	59.0	56.2
	(Standard error)	(0.7)	(1.1)	(1.5)	(1.8)	(2.0)	(2.1)	(2.6)
	Benchmark	22.9	39.2	50.8	77.9	78.3	78.6	78.8
Lost permanent job months ago	Actual	27.2	33.7	37.4	49.9	50.7	51.1	44.8
	(Standard error)	(1.4)	(2.0)	(2.7)	(2.6)	(2.7)	(3.1)	(4.2)
	Benchmark	27.2	45.3	57.4	81.3	81.5	81.6	81.7
Temp job ended months ago	Actual	27.4	35.6	42.6	65.4	65.1	63.0	65.8
	(Standard error)	(1.2)	(1.7)	(2.4)	(2.6)	(2.8)	(3.0)	(3.7)
	Benchmark	27.4	45.6	57.8	81.4	81.6	81.7	81.8
Quit months ago	Actual	17.1	21.5	28.0	41.1	44.9	41.5	38.8
	(Standard error)	(1.4)	(2.1)	(2.6)	(3.0)	(3.5)	(3.8)	(4.7)
	Benchmark	17.1	30.3	40.4	70.9	71.6	72.2	72.6
Entered labor force months ago	Actual	24.2	31.8	35.8	50.0	51.0	51.0	48.9
	(Standard error)	(0.8)	(1.1)	(1.6)	(1.6)	(1.9)	(2.1)	(2.6)
	Benchmark	24.2	41.2	53.0	79.1	79.4	79.6	79.8
Re-entered labor force months ago	Actual	16.0	22.3	25.9	35.8	37.2	37.6	34.7
	(Standard error)	(0.6)	(0.9)	(1.3)	(1.7)	(1.8)	(1.9)	(2.2)
	Benchmark	16.0	28.4	38.2	69.0	69.8	70.4	70.9

Page 1

Table 4: Subsequent Employment Probabilities by Initial Status, Actual and Benchmark, 2007: Unemployed for Months and Long-Term

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.

3.5 Changes in jobfinding rates between 2007 and 2010

The financial crisis of 2008 and the ensuing deep recession had profound effects in the labor market. Job-finding rates fell in all 16 statuses, but the composition of jobseekers shifted toward those with higher jobfinding rates.

From the peak year of 2007 to the severely depressed year of 2010, the average jobfinding rate across the 16 statuses declined from 6.3 percent to 5.8 percent. This decline of 0.5 percentage points decomposes into a component that decreased the average by 1.0 percentage points arising from lower jobfinding rates in general, and a component that increased the average by 0.7 percentage points arising from a shift of the population shares toward those with higher normal jobfinding rates. The high normal rates occur among the unemployed. The residual, a decline of 0.2 percentage points, arises from interaction effects. The tremendous change in the labor market between 2007 and 2010 left the total jobfinding flow almost unchanged, because the population shifted into unemployment, with high jobfinding rates, enough to offset the general decline of jobfinding rates across all the statuses.

A similar analysis within the unemployment statuses starts from the overall decline of 12.3 percentage points in the monthly jobfinding rate among the unemployed. Of this, 7.4 percentage points arise from declines in the rate within each status and 5.1 percentage points from a shift of the composition of unemployment toward statuses—notably loss of permanent job and long-term unemployment—with low jobfinding rates. There is also a residual of 0.2 percentage points offsetting these declines, arising from interaction effects. Within the unemployed, the shifting composition lowered jobfinding success and added to the effects of lower rates for each status.

Table 5 shows our findings for the changes in all 16 demographically adjusted jobfinding rates from 2007 to 2010. We focus on the shorter-span rates, because we are forced to omit the large job-to-job flows into employment over longer spans because of the structure of the CPS, as we discussed earlier. Recall that the short-span rates are averages over spans of one, two, and three months. Notable changes occurred in the distribution of the population among the 16 statuses: the fraction of the working-age population who were out of the labor force, wanted a job, and were available for work (*want work but not looking*) rose from 1.9 percent to 2.4 percent. The fraction working fell from 63.0 percent to 58.5 percent. Among the unemployment statuses, the furloughed fractions rose, the recently quit fraction fell, and the lost permanent job fractions rose substantially. By far the largest growth was in the long-term group, which was half a percent of the population in 2007 and 2.7 percent in 2010.

	2007			2010		
	<i>Percent of population</i>	<i>Job-finding rate</i>	<i>Contribution to total rate</i>	<i>Percent of population</i>	<i>Job-finding rate</i>	<i>Contribution to total rate</i>
Do not want work	32.2	5.4	1.74	33.0	4.4	1.47
Want work but not looking	1.87	18.2	0.34	2.41	13.8	0.33
Employed	63.0	5.0	3.17	58.5	4.4	2.58
Recently furloughed	0.20	61.9	0.12	0.23	56.6	0.13
Recently lost permanent job	0.14	41.1	0.06	0.19	30.8	0.06
Temp job recently ended	0.08	48.4	0.04	0.08	38.8	0.03
Recently quit	0.09	50.0	0.05	0.06	40.5	0.02
Recently entered labor force	0.06	27.8	0.02	0.06	18.5	0.01
Recently re-entered labor force	0.19	41.1	0.08	0.15	29.0	0.04
On furlough for months	0.22	51.0	0.11	0.32	46.2	0.15
Lost permanent job months ago	0.46	30.8	0.14	0.99	22.0	0.22
Temp job ended months ago	0.19	32.8	0.06	0.30	29.7	0.09
Quit months ago	0.20	35.2	0.07	0.19	29.0	0.05
Entered labor force months ago	0.13	22.2	0.03	0.25	14.0	0.03
Re-entered labor force months ago	0.49	30.6	0.15	0.65	23.9	0.15
Long-term unemployed	0.52	21.4	0.11	2.67	14.5	0.39
Total			6.29			5.76
Not unemployed			5.25			4.38
Unemployed			1.03			1.38

Table 5: Comparison of Short-Span Job-Finding Rates between 2007 and 2010

Job-finding rates, stated as percents of the corresponding population group who found a job, declined more or less in proportion in all statuses, in accord with the property of our model that the same index of labor-market tightness influences jobfinding rates for all types of jobseekers.

The column headed *Contribution to total rate* is the product of the population fraction in the first column and the jobfinding rate in the second column. It gives the part of the total rate, shown at the foot of the column, contributed by the people in the status corresponding to the line in the table. For example, in 2007, 32 percent of the population was in the status *do not want work*. The jobfinding rate was 5.4 percent. But this group, despite its low jobfinding rate, contributed 1.7 percentage points to the total volume of jobfinding, 6.3 percent of the working-age population each month. Workers, in the third line of the table, had the lowest jobfinding rate, 5.0 percent, but account for almost half of all jobfinding. The subtotals at the bottom of the table show that only 1.0 percentage points of the total of 6.3 percent of the population who found jobs came from the ranks of the unemployed in 2007.

A good part of the doubling of the unemployment rate that occurred between 2007 and 2010 is associated with the decline in the jobfinding rate; the rest is associated with higher flows into unemployment. In this paper, we do not measure flows into unemployment, so we do not quantify our findings in terms of unemployment rates. See Hall and Schulhofer-Wohl (2017) for a quantification using the same categories of unemployment as used here.

4 Labor-Market Tightness and the Elasticities of the Jobfinding Rates with Respect to Tightness

This section discusses our approach to measuring labor-market tightness as the average duration of vacancies. We develop an estimating equation implied by the matching function derived in Section 1, relating tightness to the jobfinding rates from the previous section. We make the case that true tightness is exogenous in the estimating equation. We then present and discuss estimates of the elasticities, using employment as an instrument to account for measurement error in observed tightness.

4.1 Data on tightness

We use data from the Job Openings and Labor Turnover Survey (JOLTS) to measure labor-market tightness, T , which is the average duration of vacancies or the ratio of the number of vacancies to the flow of hires. Figure 2 shows the number of job openings (vacancies) from JOLTS. This series traces the business cycle with high amplitude—vacancies are high in tight markets around business-cycle peaks and low in slack markets around troughs.




Figure 2: Number of Job Openings, in Thousands, from JOLTS

The flow of hires can in principle be measured from the employment side in JOLTS or from the worker side in the CPS. Figure 3 shows the number of hires from the two surveys. The CPS and JOLTS figures vary similarly over time — hires track the business cycle, but with fairly low amplitude. However, the level of hires is substantially higher in the CPS. The reasons for the discrepancy in levels 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.

JOLTS is a survey of employers and is independent of the CPS. We view the two surveys as covering overlapping but not identical subsets of the U.S. labor market. We assume that the surveys nonetheless both draw from a single labor market in the sense that a single factor, T , indexes tightness throughout the overall U.S. labor market. Under this assumption, we can calculate T as the ratio of vacancies to hires when both variables are measured consistently with each other in the same subset of the overall labor market.

The CPS has no information about vacancies, so it cannot identify tightness. This fact would remain true if we used the more standard measure of tightness as the vacancy/jobseeker ratio, usually called θ . Therefore, to measure tightness, we take both vacancies V and hires H from JOLTS. Our procedure uses the variable $T = V/H$ from JOLTS as a measure that describes the CPS labor market as well as the JOLTS labor market.

No information about jobfinding rates or matching efficiency is present in JOLTS. Under our maintained assumption that tightness is the same in the JOLTS and CPS markets, we




Figure 3: Number of Monthly Hires, in Thousands, from JOLTS and the CPS

can measure all of the objects of interest in this paper. We do not have data to test this maintained assumption.

Figure 4 shows the resulting average duration of vacancies, T . The measures of vacancies and hires are separately smoothed with annual averages before calculating T . Because vacancies vary more in proportional terms than do hires, the vacancy/hires ratio is quite procyclical.

Appendix D considers adjustments of the average duration of vacancies for recruiting intensity.

4.2 Equation for estimation of the elasticities of the jobfinding rates with respect to tightness

Our theory of matching with heterogeneous jobseekers expressed in equation (3) leads to the following equation for the measured log jobfinding rate for initial status i , over a τ -month span, in month t :

$$\log f_{i,\tau,t} = \gamma_{i,\tau,t} + \eta_i \log T_t + \epsilon_{i,\tau,t}, \quad (7)$$

where $f_{i,\tau,t}$ is the observed jobfinding rate; $\gamma_{i,\tau,t}$ contains constants and trends; and $\epsilon_{i,\tau,t}$ is the detrended matching efficiency of type i at date t over the subsequent span τ , the residual in the equation. The parameter η_i is the elasticity of jobfinding with respect to tightness, T_t . We assume that

$$\gamma_{i,\tau,t} = \alpha_{i,\tau} + \delta_{i,\tau} t + \omega_{i,\tau} t^+ + \psi_{i,\tau,s(t)}, \quad (8)$$




Figure 4: Average Duration of Vacancies, Calculated from JOLTS

where $\delta_{i,\tau}t$ is a linear trend that operates over the whole sample, $\omega_{i,\tau}t^+$ is an additional trend starting in January 2008, $s(t) = 1, \dots, 12$ is the month of observation t , and $\psi_{i,\tau,s(t)}$ is a seasonal effect.

This model has a separate equation for each $\{i, \tau\}$ pair—there are no cross-equation restrictions. We can therefore estimate it separately for each initial status i and span τ . Although this approach is not minimum variance if the disturbances are correlated across statuses, we find that the correlations across the estimated residuals are small. Also, given the large number of estimated coefficients relative to the number of data points, we do not believe that a generalized least squares estimation procedure would be appropriate.

We use the estimates of jobfinding rates adjusted for the changing characteristics of the population, as discussed earlier, as the left-hand variable of equation (7). We doubt that combining the two estimation stages would be practical and have no reason to believe it would affect our conclusions. Our bootstrap standard errors take both stages into account.

In estimating the equation, we average the three short spans (one, two, and three months after the conditioning status) to form the jobfinding rate for the first span category, called *short*, and the four longer spans (12 through 15 months) to form the second jobfinding rate category, called *long*. We estimate the equation on monthly observations from January 2001 through December 2013.

4.3 Simultaneity and identification

The jobfinding rate f and labor-market tightness T are jointly determined endogenous variables. Within the broad class of labor-market models associated with Diamond, Mortensen, and Pissarides, the two variables are determined in a two-equation system,

$$\log f = \alpha + \eta \log T + \epsilon \tag{9}$$

and

$$\chi T = J. \tag{10}$$

For simplicity, we consider the case of a single type of jobseeker. The second equation expresses the zero-profit condition— χ is the monthly cost of maintaining a vacancy, so χT is the expected cost of hiring one new worker, and that amount equals the payoff J to making a new hire.

A key issue of identification is what happens to tightness T if matching efficiency ϵ changes. From equation (10), tightness can only change when ϵ changes, to the extent that J changes. J is the present value over the duration of the job of the difference between the worker’s marginal product and the wage. An increase in ϵ will presumably raise the jobfinding rate, which will lower the unemployment rate, raise employment, and lower the marginal product of labor. Chodorow-Reich and Karabarbounis (2016) show that productivity changes carry over to opportunity cost changes, so there is no effect on the difference between productivity and the wage as productivity varies. Thus we take the effect on T operating through the marginal product of labor to be zero. The increase in the jobfinding rate will raise the value of the worker’s outside option in the wage bargain—to the extent that the outside option is influential in the bargain, tightness will fall. But Hall (2017) presents evidence that the outside option has almost no effect on the wage bargain, in a realistic alternating-offer setting. The influence of the outside option under the conventional assumption of a Nash wage bargain arises from the unrealistic influence of the irrelevant option to discontinue bargaining that is implicit in the Nash bargain setup.

Our assumption that shocks to matching efficiency do not affect our measure of tightness, T , implies that the DMP measure of tightness in the labor market rises with matching efficiency, a standard conclusion in the DMP literature—see Appendix A.

In principle, changes in ϵ could induce changes in the discount rate that would influence T , but we do not believe that these would be important. Hall (2017) shows that it takes quite large changes in discounts to change tightness materially.

Our overall conclusion is that spontaneous movements in ϵ have essentially full direct effects on $\log f$ and small effects on T . Having excluded fluctuations in T from certain channels, we need to explain the sources of the high observed volatility of T . That is, we need to indicate what causes the job value J to fluctuate. Hall (2017) shows that large

measured fluctuations in financial discounts results in volatility of J , a financial present value. Another source could be spontaneous volatility of the opportunity cost of work, which would move the wage relative to the marginal product of labor and generate movements of J . In the presence of product market power, J is the present value of the difference between the marginal revenue product of labor and the wage. Thus a third source of volatility in T could be fluctuations in markups, which create changes in the marginal revenue product—see Rotemberg and Woodford (1999).

Appendix E presents evidence that our measure of T contains an important component that appears to reflect high-frequency measurement error. In other words, despite our belief that there is no fundamental feedback from the disturbance ϵ to the level of true tightness, there is probably an errors-in-variables problem in estimation. This finding supports the use of an instrumental variables (two-stage least squares) estimator. As the instrumental variable, we use the employment count from the Bureau of Labor Statistics’ Current Employment Statistics payroll survey, which has a large number of responding firms and correspondingly low sampling errors. Our identifying hypothesis is that the much larger payroll survey is free of the short-term noise arising from the sampling error in JOLTS. We form moving averages of the JOLTS measure of T as the endogenous variable and similar moving averages of the employment count as the instrument—these are equally weighted over the current and the future 2 or 14 months for short and long spans, respectively.

4.4 Estimates

Our objective is to estimate the parameters η_i , the elasticities of the jobfinding rates of jobseekers with respect to tightness T , across status i . To help interpret the estimated elasticities, recall that, in the case of a single type of jobseeker, the elasticity of the matching function is $\eta/(1 + \eta)$. For example, an estimate of $\eta = 1$ corresponds, in this sense, to a matching elasticity of 0.5.

Table 6 shows estimates of the based on equation (7), using data for the JOLTS period, 2001 through 2013. The left panel refers to short spans and the right panel to long spans. In each panel, the left column is the estimated elasticity, with standard error below. The middle column shows the trend for the entire period since 2001, in percent per year. The right column shows the extra trend starting in 2008.

All of the elasticities of the jobfinding rate with respect to tightness are positive and most have small bootstrap standard errors. The elasticities have substantial heterogeneity—the evidence against equal elasticities (the case of a Cobb-Douglas overall matching function) is quite strong. For both short and long spans, the *recently lost permanent job* and *lost permanent job months ago* initial statuses have high elasticities of jobfinding with respect to tightness.

<i>Initial status</i>	<i>Short span</i>			<i>Long span</i>		
	<i>Elasticity with respect to vacancy duration</i>	<i>Trend in efficiency, 2001-2013, percent per year</i>	<i>Additional trend in efficiency, 2008-2013, percent per year</i>	<i>Elasticity with respect to vacancy duration</i>	<i>Trend in efficiency, 2001-2013, percent per year</i>	<i>Additional trend in efficiency, 2008-2013, percent per year</i>
Do not want work (Standard error)	0.50 (0.04)	-2.54 (0.20)	-1.98 (0.37)	0.60 (0.05)	-2.65 (0.22)	-1.06 (0.42)
Want work but not looking (Standard error)	0.79 (0.07)	-2.46 (0.33)	-2.06 (0.61)	0.76 (0.08)	-2.90 (0.38)	0.70 (0.70)
Employed (Standard error)	0.35 (0.03)	-2.78 (0.14)	0.19 (0.28)			
Recently furloughed (Standard error)	0.21 (0.07)	0.11 (0.33)	-1.22 (0.71)	0.34 (0.11)	-2.09 (0.48)	3.58 (0.91)
Recently lost permanent job (Standard error)	0.96 (0.16)	-1.05 (0.62)	-2.44 (1.32)	0.77 (0.15)	-3.06 (0.68)	1.50 (1.25)
Temp job recently ended (Standard error)	0.47 (0.17)	-0.87 (0.71)	-1.84 (1.54)	0.70 (0.26)	-3.56 (1.07)	2.78 (2.00)
Recently quit (Standard error)	0.48 (0.19)	-0.51 (0.67)	-3.58 (1.41)	0.47 (0.25)	-2.24 (0.95)	2.18 (1.82)
Recently entered labor force (Standard error)	0.77 (0.40)	-3.06 (1.76)	-4.86 (3.38)	0.13 (0.58)	-0.25 (2.39)	-0.54 (4.25)
Recently re-entered labor force (Standard error)	0.76 (0.14)	-2.46 (0.61)	-1.97 (1.22)	0.65 (0.18)	-1.59 (0.70)	-0.69 (1.36)
On furlough for months (Standard error)	0.50 (0.10)	-0.54 (0.55)	-1.20 (1.14)	0.38 (0.15)	-2.29 (0.55)	2.82 (1.23)
Lost permanent job months ago (Standard error)	1.21 (0.11)	-2.16 (0.50)	-2.12 (1.15)	0.93 (0.12)	-4.57 (0.44)	3.50 (0.84)
Temp job ended months ago (Standard error)	0.75 (0.16)	-1.64 (0.71)	-1.48 (1.43)	0.46 (0.19)	-3.06 (0.85)	2.89 (1.68)
Quit months ago (Standard error)	0.83 (0.17)	-2.56 (0.68)	-0.42 (1.45)	0.98 (0.20)	-2.24 (0.83)	-1.35 (1.51)
Entered labor force months ago (Standard error)	1.24 (0.24)	-3.37 (1.41)	-3.24 (2.53)	0.78 (0.28)	-2.30 (1.68)	2.45 (2.84)
Re-entered labor force months ag (Standard error)	0.83 (0.11)	-2.45 (0.47)	-1.83 (1.01)	0.69 (0.12)	-2.43 (0.45)	-0.11 (0.94)
Long-term unemployed (Standard error)	1.23 (0.11)	-3.89 (0.64)	-2.19 (1.12)	0.51 (0.12)	-3.75 (0.65)	3.55 (1.25)

Table 6: Elasticity and Trend Estimates

The composition of jobseekers changes in recessions for two reasons. First, the mix of reasons why people leave jobs can change. Second, our method reveals a dynamic effect that occurs because recessions do not affect all jobseekers equally. We find that jobfinding rates are much more responsive to labor market tightness for some categories of jobseekers, such as losers of permanent jobs, than for other categories, such as people who were recently laid off. When a recession hits and tightness falls, the categories with higher elasticities experience larger reductions in jobfinding rates. As a result, these categories grow to make up a larger share of the pool of jobseekers. We show that accounting for these compositional changes is fundamental to proper measurement of matching efficiency.

Without the sole exception of the short-span rate for the *recently laid off* category, jobfinding rates adjusted for changes in labor-market tightness trended downward over the period from 2001 through 2007. The downward trend is particularly steep for the long-term unemployed (more than six months). The downward trend in short-span matching efficiencies generally declined faster in 2008 and later. On the other hand, in the majority of the initial-status categories, the earlier downward trend in efficiency for long spans reversed partially or even fully in the period starting in 2008. There is little support for the hypothesis that matching efficiency fell after the financial crisis and the following long period of turmoil.

The average across the short-span elasticities is 0.527, with bootstrap standard error 0.060. The corresponding elasticity of the matching function with respect to vacancies, if all the elasticities had this value, from equation (5), is 0.345, and the elasticity with respect to equal proportional increases in all statuses, is one minus this amount, 0.655. Both have bootstrap standard errors of 0.026. For the short-span equations, the matching elasticity estimate 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 Matching Efficiency

We use the residuals and estimated trends from equation (7) to compute indexes of matching efficiency. The residuals,


$$\epsilon_{i,t,\tau} = \log f_{i,\tau,t} - [\alpha_{i,\tau} + \delta_{i,\tau}t + \omega_{i,\tau}t^+ + \psi_{i,\tau,s(t)} + \eta_i \log T_t], \quad (11)$$

measure the detrended matching efficiency of type i over span τ , or, equivalently, the deviation of the observed job-finding rate from its status- and span-specific constant level and trend, adjusted for changes in labor-market tightness. We can also measure the trend in matching efficiency by adding back our estimates of $\delta_{i,\tau}t + \omega_{i,\tau}t^+$ to the residuals. The residuals also include measurement error in jobfinding rates, but such measurement errors should average to zero over time. In particular, our presentation of the results focuses mainly on


annual averages, so much of the measurement error should average out over the course of each year.

5.1 Indexes of matching efficiency calculated from our estimates


We calculate indexes of matching efficiency for each of the 16 initial labor-market statuses over short spans, and for all initial statuses except employment over long spans. Because we hold the distribution of individuals' characteristics constant in calculating the jobfinding rates on the left-hand side of equation (7), the movements in these indexes are insulated from changes in the distribution of characteristics. Figure 5 shows the resulting detrended indexes for 9 of the more important statuses. These are the exponentials of the values described in equation (11) and are indexes normalized to one in 2001. The trends are shown in Table 6 and allow for different trends, generally downward, in matching efficiency in the pre-crisis period 2001 through 2007 and in the following period, 2008 through 2013.




(a) Not in LF




(b) Want job




(c) On layoff for months




(d) Lost permanent job months ago




(e) Temp job ended months ago




(f) Quit months ago



(g) Entered months ago



(h) Re-entered months ago



(i) Long-term unemployed


Figure 5: Detrended Matching Efficiency for Nine Statuses (Short Spans in Blue and Long Spans in Orange)

The overall impression from the 9 categories shown in Figure 5, and confirmed for the remaining 7 categories not shown, is that movements in matching efficiency around its downward trend since 2001 are generally small and unsystematic. To put it differently, the tightness measure is able to take account of changes in the labor market when estimation occurs over the relatively mild recession of 2001 and the deep recession that started at the end of 2007, after adjustment for the two-part trend in our setup. We noted earlier that the correlation across the 16 categories of the residual—the measure of matching efficiency—is essentially zero. The only respect in which our estimation procedure lowers the correlations is the common set of right-hand variables in the equations: log tightness, seasonal effects, and the two-part trend. The absence of a non-trend, non-seasonal, non-tightness common element in the residuals is evidence in favor of our specification.


The pattern of annual matching efficiency for the initial status *lost permanent job months ago* is representative in terms of its movement over time, and more precisely estimated because large numbers of jobseekers fell into this category. For both short and long spans, Table 6 shows downward trends in 2001 through 2007, and almost no trend in 2008 through 2013. Both measures of detrended efficiency rose during the recovery from the 2001 recession, and fell as the economy reached its peak in 2007. After 2008, long-span efficiency remain constant, while short-span efficiency rose at the end of our sample period.

In the closely watched category *long-term unemployed*, where the estimated trend over 2008 through 2013 was also close to zero, the short- and long-span indexes move generally together, rising in the recovery following the 2001 recession, falling from 2004 to the peak in 2007, flattening during the recession, then growing at the very end.

Figure 6 shows the indexes without subtraction of the trend terms in equation (7). The scale is compressed relative to the detrended indexes in Figure 5 to accommodate the trends. Notice that the trends are downward over time for all of the initial statuses shown, corresponding to the ratios of 2013 jobfinding rates to 2003 rates in Table 2 that are almost all below one. In most of the categories, the downward trend in efficiency after 2008 is less than in the earlier period, for the long-span measure. In particular, despite the huge increase in unemployment after 2007, there is little sign of any corresponding movement of matching efficiency.




(a) Not in LF




(b) Want job




(c) On layoff for months




(d) Lost permanent job months ago




(e) Temp job ended months ago




(f) Quit months ago



(g) Entered months ago




(h) Re-entered months ago




(i) Long-term unemployed

Figure 6: Matching Efficiency for Nine Statuses, Including Trend (Short Spans in Blue and Long Spans in Orange)



(a) Detrended



(b) With trend

Figure 7: Weighted Average Matching Efficiency Trend

The left side of Figure 7 shows indexes of weighted average matching efficiency including removal of trends. The indexes use weights calculated as the shares of the components in the population in the three years preceding the crisis, 2005 through 2007. Because the jobfinding rates underlying the indexes hold constant the distribution of worker characteristics conditional on labor-market status, the aggregate indexes hold 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 jobseekers among the initial statuses. The index for short spans includes job-to-job movers while the one for long spans includes only the unemployed and people not in the labor force. Long-span efficiency moves much the same way as short-span. Both are quite smooth—adding over the 16 categories smooths away most of the volatility shown in Figure 5 and Figure 6. The right side of Figure 7 shows the same data without adjustment for trend.

Appendix C discusses three alternative specifications in terms of their implications for the indexes of overall matching efficiency. The first is the same as our baseline specification except that no demographic effects are swept out in a preliminary estimation. The second is similar to the base except that the elasticity of the jobfinding rate with respect to tightness is constrained to be the same for every initial status group. The third is the same as the base specification, but uses only data for 2001 through 2007, the years prior to the crisis. Our basic conclusion holds in all three alternatives that proper accounting for heterogeneity among jobseekers results in an index of matching efficiency that follows a smooth trend with no special movement in the years after the crisis in 2008. We also constructed Divisia-style

indexes with time-varying weights. The difference between these indexes and our fixed-weight indexes was tiny.

5.2 Measuring matching efficiency when there is only one type of jobseeker

Suppose that there is only one type of jobseeker, an unemployed person, without regard to the type of unemployment. We explore this approach because much of the literature on the matching function takes the count of unemployed jobseekers as the single job-seeking input to the function.

The relation among the jobfinding rate f_t , matching efficiency, γ_t , and tightness, T_t , is

$$\log f_t = \log \gamma_t + \eta \log T_t. \quad (12)$$

The jobfinding rate is the ratio of hires out of unemployment, H_t^U , to the number of unemployed people, U_t . Thus single-type matching efficiency is

$$\gamma_t = \frac{H_t^U / U_t}{T_t^\eta}. \quad (13)$$

To appraise this measure, we take H_t^U from the BLS’s annual series on flows from unemployment to employment in the CPS. We use an elasticity of $\eta = 0.74$ —the average of the estimates of the elasticity in Table 6. Figure 8 compares our measure of matching efficiency to the measure that uses unemployment as the sole measure of job-seeking volume. The results show that the approach using only total unemployment is misleading relative to one that includes all types of jobseekers and that recognizes heterogeneity among the types. The single-type equation finds large movements in γ_t that arise from changes in the composition of the unemployed and not from shifts in matching efficiency for individual types.

The single-type measure considerably overstates the decline in matching efficiency between 2007 and 2010, the period when unemployment doubled. It infers a substantial decline in efficiency from its measure of the jobfinding rate, H_t^U / U_t . But this measure overstates the decline in the rate because its denominator does not take account of the shift of the composition of unemployment toward those with low job-finding rates. The denominator overstates the rise in effective jobseeking among the unemployed by treating them as homogeneous. In addition, the single-type measure does not consider jobseeking success among the large number of jobseekers who are not counted as unemployed. These biases drove the ratio down and created the illusion of a major decline in matching efficiency, when, by our mix-adjusted comprehensive measure, almost no decline occurred relative to the normal downward trend. The same distortion operated in the recession of 2001 and its aftermath, though not nearly as dramatically.




Figure 8: Comparison of this Paper’s Measure of Matching Efficiency to a Naive Measure

6 Related Research

Mismatch has been an important topic of research related to this paper. Barnichon and Figura (2015b) incorporate an explicit term in their equation for mismatch, resulting from the curvature of the matching function. Because the term makes little contribution, and for the reasons laid out in Appendix E, we do not include that term, though we agree that mismatch is a phenomenon of the labor market that is quantitatively important in other contexts. Şahin, Song, Topa and Violante (2014) measure mismatch as the deviation of actual unemployment from the planner’s allocation of searchers to markets. They find that mismatch across industries and occupations accounts for somewhat under 20 percent of the increase in unemployment during the Great Recession, while geographic mismatch is insignificant. Their findings support the hypothesis that the dominant source of movements in tightness is an aggregate influence acting on all sectors.

A large literature, surveyed recently in Elsby, Michaels and Ratner (2015a), builds models of search intensity. Variation in intensity is potentially an important determinant of what we measure. We do not break down matching efficiency into components of search effort and jobfinding success per unit of search effort, because we do not measure search effort directly; absent direct measures of search effort, as in Krueger and Mueller (2011), we cannot break the two factors apart. (See Hornstein and Kudlyak (2015) for a setup where endogenous search effort is inferred indirectly.) Because we use jobfinding rates that hold demographics constant, our approach to estimation does make adjustments for differences in search effort

associated with observed demographics. Still, 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. We also do not consider the distinction between a contact of a jobseeker and employer and the creation of a job match. The probability that a contact results in a hire is one of the factors determining the jobfinding rates that we measure.

We treat vacancies as homogeneous in our base specification. In principle, vacancies should be disaggregated to recognize their heterogeneity and likely variations in recruiting intensity. Davis, Faberman and Haltiwanger (2013) and Gavazza, Mongey and Violante (2016) are important recent studies of that heterogeneity. In Appendix D, we show that incorporation of recruiting intensity as modeled in those papers has essentially no effect on our conclusions, though it has some effect on the trend rate of decline in efficiency.

Many discussions of the matching process in the labor market are organized around the Beveridge curve, which portrays movements of unemployment and job vacancies. Shifts in matching efficiency are one source of instability in the Beveridge curve. Changes in the inflow rate to unemployment are another. The large changes in the composition of unemployment over the business cycle are major sources of shifts. This paper focuses only on matching efficiency and not on other shifters of the curve, so we do not try to express our findings in terms of the Beveridge curve. Our finding of stability of matching efficiency at the level of different types of jobseekers is consistent with large shifts in the curve arising from those other sources. Because the Beveridge curve concerns unemployment and not the other important sources of jobfinding, the Beveridge-curve framework does not provide a comprehensive view of flows into employment.

References

- Abowd, John M. and Arnold Zellner, “Estimating Gross Labor-Force Flows,” *Journal of Business & Economic Statistics*, 1985, 3 (3), 254–283.
- Ahn, Hie Joo and James D. Hamilton, “Heterogeneity and Unemployment Dynamics,” August 2016. Department of Economics, University of California, San Diego.
- Barnichon, Regis and Andrew Figura, “Declining Desire to Work and Downward Trends in Unemployment and Participation,” *NBER Macroeconomics Annual*, 2015, pp. 449–494.
- and —, “Labor Market Heterogeneity and the Aggregate Matching Function,” *American Economic Journal: Macroeconomics*, October 2015, 7 (4), 222–49.
- 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.
- Chodorow-Reich, Gabriel and Loukas Karabarbounis, “The Cyclicity of the Opportunity Cost of Employment,” *Journal of Political Economy*, 2016. Forthcoming.

- Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante, “Mismatch Unemployment,” *American Economic Review*, November 2014, 104 (11), 3529–64.
- Davis, Steven J. and John Haltiwanger, “Labor Market Fluidity and Economic Performance,” *Proceedings of the Jackson Hole Symposium, Federal Reserve Bank of Kansas*, September 2014, pp. 17–107.
- , 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., Ryan Michaels, and David Ratner, “The Beveridge Curve: A Survey,” *Journal of Economic Literature*, 2015, 53 (3), 571–630.
- Elsby, Michael W.L., Bart Hobijn, and Ayşegül Şahin, “On the Importance of the Participation Margin for Labor Market Fluctuations,” *Journal of Monetary Economics*, 2015, 72, 64 – 82.
- , —, —, 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.
- Feng, Shuaizhang and Yingyao Hu, “Misclassification Errors and the Underestimation of the US Unemployment Rate,” *American Economic Review*, 2013, 103 (2), 1054–70.
- Fujita, Shigeru and Giuseppe Moscarini, “Recall and Unemployment,” Working Paper 19640, National Bureau of Economic Research, November 2013.
- Gavazza, Alessandro, Simon Mongey, and Giovanni L Violante, “Aggregate Recruiting Intensity,” Technical Report, National Bureau of Economic Research 2016.
- Hall, Robert E., “High Discounts and High Unemployment,” *American Economic Review*, February 2017, 107 (2), 305–330.
- and Sam Schulhofer-Wohl, “The Pervasive Importance of Tightness in Labor-Market Volatility,” April 2017. Hoover Institution, Stanford University.
- Hornstein, Andreas and Marianna Kudlyak, “Estimating Matching Effectciency with Variable Search Effort,” May 2015. Federal Reserve Bank of Richmond.
- Hyatt, Henry R. and James R. Spletzer, “The Recent Decline in Employment Dynamics,” March 2013. Center for Economic Studies, US Census Bureau.
- Katz, Lawrence F., “Layoffs, Recall, and the Duration of Unemployment,” Technical Report, National Bureau of Economic Research, 1986. Working Paper 1825,.
- 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.
- 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.
- Papke, Leslie E. and Jeffrey M. Wooldridge, "Econometric methods for fractional response variables with an application to 401(k) plan participation rates," *Journal of Applied Econometrics*, 1996, 11 (6), 619–632.
- 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.
- Rotemberg, Julio J. and Michael Woodford, "The Cyclical Behavior of Prices and Costs," 1999. in Taylor and Woodford, eds., *Handbook of Macroeconomics*, Chapter 16, pp. 1051–1135, 1999.
- Veracierto, Marcelo, "Worker flows and matching efficiency," *Economic Perspectives*, 4th quarter 2011, pp. 147–169.

Appendixes for Online Publication

A Relation between the Standard DMP Matching Setup and the One in this Paper

A.1 The standard DMP setup

Start from the matching function:

$$H = \mu P^\nu V^{1-\nu}. \quad (14)$$

Tightness:

$$\theta = \frac{V}{P}. \quad (15)$$

Job-finding rate:

$$f = \frac{H}{P} = \mu \theta^{1-\nu}. \quad (16)$$

Job-filling rate:

$$q = \frac{H}{V} = \mu \theta^{-\nu}. \quad (17)$$

Zero-profit condition:

$$\frac{\kappa}{q} = J. \quad (18)$$

Vacancies:

$$V = \theta P. \quad (19)$$

A.2 The paper's setup with one type of jobseeker

Tightness (vacancy duration):

$$T = \frac{V}{H}. \quad (20)$$

Start from the job-finding rate:

$$f = \gamma T^\eta. \quad (21)$$

Matching function:

$$H = \gamma T^\eta P; H = \gamma^{1/(1+\eta)} V^{\eta/(1+\eta)} P^{1/(1+\eta)}. \quad (22)$$

Job-filling rate:

$$\frac{H}{V} = \frac{1}{T}. \quad (23)$$

Zero-profit condition:

$$\kappa T = J \quad (24)$$

Vacancies:

$$V = T \cdot H. \quad (25)$$

A.3 Implications of holding T constant as μ changes

Because $T = 1/q$, holding T constant is equivalent to holding q constant. From equation (17), an increase in μ must be accompanied by an increase in DMP tightness, θ , to hold q constant.

A.4 The paper's setup with multiple types of jobseekers

Tightness (vacancy duration):

$$T = \frac{V}{H}. \quad (26)$$

Start from the job-finding rates:

$$f_i = \gamma_i T^{\eta_i}. \quad (27)$$

Matching function satisfies:

$$H = \sum_i P_i \gamma_i \left(\frac{V}{H} \right)^{\eta_i}. \quad (28)$$

Job-filling rate:

$$\frac{H}{V} = \frac{1}{T}. \quad (29)$$

Zero-profit condition:

$$\kappa T = J. \quad (30)$$

Vacancies:

$$V = T \cdot H. \quad (31)$$

The property that the zero-profit condition involves only T is important for our approach to identification.

B Attrition in the CPS

Table 7 describes our success in matching respondents in different months in the CPS. It shows the weighted percent of observations that were successfully matched to an observation on the same person some month later, conditional on the initial observation being early enough in the CPS sample rotation that a match was theoretically possible. (For example, a match one month later is theoretically possible if the initial observation is not in the outgoing rotation group; a match 15 months later is theoretically possible only if the initial observation is in the incoming rotation group.) The intervals correspond to the spans that we use for estimation. The short-span match rates are quite high; the long-span match rates less so. We calculated the success rates by year. The bottom line of the upper panel shows the standard deviations of the rates across years. They are uniformly small; the success rates were stable over the period from 2001 through 2013.

The table also shows the matching rates that we would obtain if we used the method of Madrian and Lefgren (2000). That method produces a slightly higher match rate than Nekarda's method at a horizon of 1 month because the Madrian and Lefgren 1-month match

Number of months separating observations	1	2	3	12	13	14	15
Percent matched, Nekarda method	93.6	91.3	89.3	75.3	74.5	73.5	72.5
Standard deviation across years	0.2	0.3	0.3	1.9	1.9	2.0	2.1
Percent matched, Lefgren-Madrian method	94.7	90.3	86.3	68.8	66.4	64.2	62.1
Standard deviation across years	0.2	0.3	0.5	1.9	1.9	1.9	1.9

Table 7: Percent of Observations Matched between Months in the Current Population Survey

does not condition on what happens in subsequent months, while Nekarda’s 1-month match does. However, the Madrian and Lefgren match rates are lower than the Nekarda match at all horizons longer than 1 month, because at longer horizons Nekarda’s method allows some matches that Madrian and Lefgren’s method rejects.

Following standard principles of attrition adjustment, we offset the potential bias caused by higher weighting of the respondents who are less likely to drop out. For each date t and span τ , we estimate a fractional logit model (Papke and Wooldridge (1996)) for the probability that an individual observed at t is also observed at $t + \tau$, as a function of the same variables that are on the right-hand side of our logit for jobfinding rates. Let $\hat{p}_{i,t,\tau}$ be the predicted probabilities of remaining in the sample from this model for individual i observed at t , over a span of τ months. To estimate the jobfinding rates over a span of τ months from the logit equation, we weight each observation by $1/\hat{p}_{i,t,\tau}$ times the product of Nekarda’s linking weight and the survey weight. Thus observations with a lower probability of remaining in the sample are given higher weight. We re-estimate the weights for each bootstrap sample. We use a fractional logit model because *remaining in the sample* is not a binary event with Nekarda’s weights and so cannot be the dependent variable in a conventional logit model.

Reweighting to account for attrition did not change the estimated jobfinding rates appreciably. This finding is unsurprising because the variables in the attrition model are also controls in the model for jobfinding rates. In essence, the attrition weights account only for potential misspecification of the functional form of the jobfinding rate equation.

C Estimates of Aggregate Matching Efficiency for Alternative Specifications

Table 8 shows our basic results for three alternative specifications. The left panel shows the detrended index of matching efficiency measured over short spans and the right panel the index for long spans, excluding job-to-job, as in Figure 7. It includes the years 2001 through 2003, years affected by the 2001 tech crash, and 2008, years affected by the financial crisis. The left column in each panel repeats the index from the body of this paper. The next column is similar in all respects except that no demographic effects are swept out. The third

<i>Short spans</i>					<i>Long spans</i>			
<i>Year</i>	<i>Base</i>	<i>No demo-graphics</i>	<i>Common elasticity</i>	<i>Pre-crisis</i>	<i>Base</i>	<i>No demo-graphics</i>	<i>Common elasticity</i>	<i>Pre-crisis</i>
2001	1.000	1.000	0.040	1.000	1.000	1.000	1.000	1.000
2002	1.033	1.035	1.076	0.923	1.028	1.034	1.029	0.972
2003	1.030	1.032	1.083	0.875	1.068	1.089	1.068	0.989
2008	1.005	1.005	1.022		1.020	1.015	1.020	
2009	1.039	1.037	1.133		1.044	1.046	1.043	
2010	1.026	1.029	1.086		1.029	1.051	1.029	
Standard deviation	0.024	0.024	0.040	0.029	0.023	0.027	0.023	0.020

Note: The standard deviation includes the omitted years 2004-2007 and 2011-2013

Table 8: Detrended Indexes of Matching Efficiency for Alternative Specifications

column is similar to the base except that the elasticity of the jobfinding rate with respect to tightness is constrained to be the same for every initial status group. This corresponds to the assumption that the matching function is Cobb-Douglas in a weighted sum of jobseekers. The right-most column is based on estimates of the base specification, but uses only data for 2001 through 2007, the years prior to the crisis.

In all cases, the results conform to the overall conclusion of the paper, that a fixed-weight index shows that matching efficiency departed from its trend only slightly. The standard deviations of the alternative indexes of matching efficiency, shown at the foot of the table, are greater than the preferred base specification, shown at the left of each panel, but are still quite small. By contrast, the standard deviation of the detrended version of the matching efficiency index in Figure 8, based on a single type of unemployment, is vastly higher, at 0.155.

D Recruiting Intensity

D.1 Davis-Faberman-Haltiwanger's Estimates of Vacancy Duration and Recruiting Intensity

Davis et al. (2013) derive an adjustment to the JOLTS measure of vacancy duration to account for time aggregation, and a second adjustment for recruiting intensity. Table 9 shows, in a format similar to Table 8, our measure of matching efficiency with trend, using T with the time-aggregation adjustment alone, and T with the product of the time-aggregation adjustment and the recruiting-efficiency adjustment. Adding the time-aggregation adjustment by itself has almost no effect on the index of matching efficiency. The adjustment for recruiting intensity eliminates some of the downward trend in both indexes, but does not affect the overall conclusion of the paper.

<i>Short spans, with trend</i>				<i>Long spans, with trend</i>		
<i>Year</i>	<i>Base</i>	<i>With time aggregation</i>	<i>With time aggregation and intensity</i>	<i>Base</i>	<i>With time aggregation</i>	<i>With time aggregation and intensity</i>
2001	1.000	1.000	1.000	1.000	1.000	1.000
2002	1.006	1.011	1.004	1.000	1.004	0.999
2003	0.978	0.984	0.980	1.012	1.020	1.007
2008	0.833	0.835	0.869	0.843	0.845	0.893
2009	0.833	0.837	0.869	0.833	0.842	0.882
2010	0.796	0.804	0.832	0.795	0.803	0.849
2013	0.743	0.751	0.786	0.743	0.753	0.807

Table 9: Indexes of Matching Efficiency, with Trends, Including Adjustments from Davis and Co-Authors

D.2 Implications of the Findings about Recruiting Intensity of Gavazza and Co-Authors

Gavazza et al. (2016) conclude that, to a fair approximation, recruiting intensity satisfies their equation (29):

$$\log \Phi = -\alpha \frac{\gamma_2}{\gamma_1 + \gamma_2} \log q = \pi \log T, \quad (32)$$

and the value of the elasticity is $\pi = 0.40$, from the values of the parameters reported in their Table 2. If the matching function is written with endogenous recruiting intensity that enters the matching function in vacancy-augmenting form, as in their equation (2), it is

$$H = \sum_i \phi_i \left(\frac{\Phi V}{H} \right) P_i = \sum_i \gamma_i T^\pi T^{\tilde{\eta}_i}. \quad (33)$$

Thus we can partition our estimated elasticities η_i into lower values $\tilde{\eta}_i = \eta_i - \pi$ that isolate the effect from tightness itself and the common effect of endogenous recruiting intensity, π . The addition of endogenous recruiting effort requires no modification in our estimation. It only contributes this decomposition of our findings.

E Mismatch Effects in the Duration of Vacancies, T

Our estimation equation (7) involves a concave log transform of T , so there is an issue of using the aggregate value when there is dispersion across units— T is potentially subject to mismatch bias. To understand this issue, we studied the industry-level data on hires and job openings published for JOLTS, across 4 geographic units and 25 industries. We hypothesize that the ratio of openings to hires—the vacancy duration T —has three components: (1) an aggregate tightness measure, as used in our work, (2) a unit-specific component, reflecting the deviation of tightness in the unit from the aggregate measure, and (3) measurement error,

Lag, months	12	18	24
Coefficient	0.226	0.087	0.036
Standard error	(0.015)	(0.015)	(0.015)

Table 10: Estimates for the Forecasting Power of Lagged Vacancy Duration at Selected Lags

occurring because JOLTS is a fairly small survey and from other random sources unrelated to tightness.

We believe that that the unit-specific component is moderately persistent, mainly because both aggregate vacancy duration and its counterpart in the geographic and industry units are persistent. On the other hand, the measurement errors are likely to be transitory. We begin our investigation by studying the autocovariance functions of the disaggregated data stated as deviations from the aggregate series for T_t . If the measurement errors were white noise and the unit-specific tightness process quite persistent, the functions would spike at zero—the only lag value where the measurement error would contribute—then drop immediately to a gradually declining value starting at a lag of one month. In fact, the autocovariance functions resemble those of a fairly non-persistent autoregressive process, with no special spike at zero lag. There is no highly persistent component.

To capture this finding more rigorously, we estimate the following equation:

$$\log T_{i,t} = \lambda_i + \bar{T}_t + \rho_\ell \log T_{i,t-\ell} + \nu_{i,t}. \quad (34)$$

Here λ_i is a level effect for industry i , \bar{T}_t is a time effect, ρ_ℓ measures the predictive power of the observation ℓ months earlier, and $\nu_{i,t}$ is the residual. The coefficient ρ_ℓ declines with the length of the lag, ℓ —it is analogous to the autocovariance in this panel setting.

Table 10 shows estimates of the prediction coefficient.

F Technical Issues in Computing Jobfinding Rates in Small Cells

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 jobfinding rate to be 1 or 0.

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 (6) for that initial status, date, and time span. Instead, we impute the 15-month jobfinding rates for new entrants in February 2008 based on the jobfinding 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 jobfinding rates when a particular bootstrap sample has no observations for a given initial status, date, and time span.