1
2016
Economic perspectives

What does online job search tell us about the labor market?

R. Jason Faberman and Marianna Kudlyak

Introduction and summary

Online job search (OJS) has become a pervasive part of job-finding and hiring in the U.S. labor market. Its use by both job seekers and employers has exploded in a fairly short period. There are now a variety of well-known, high-traffic websites devoted to job search. Furthermore, online job search has become the preferred method of search for nearly all types of job seekers. Recent research, as well as new findings that we present in this article, suggests that online job search is the new norm for how job seekers find work and how employers hire workers.

The microdata on online job search enable new and exciting avenues for labor research that were not possible with existing survey data. In particular, online job search data often allow researchers to follow job seekers' daily search behavior throughout their search spells. This provides researchers with the opportunity to study the entire search process and factors that influence it. This process is an important area of modern labor and macroeconomics research. Specifically, contemporary models of labor market search characterize job search and hiring through the use of a matching function—which is often "black box," or opaque, in nature. Online job search data provide an opportunity for researchers to help open up the black box.

Until even a few years ago, online job search had not been studied very much in economic research. Yet there were some notable exceptions. One key early study on OJS was by Kuhn and Skuterud (2004). Their work found that in 2000, OJS was used by about one-quarter of unemployed job seekers and was no more effective than traditional job search methods in helping them find work. More-recent research, however, suggests that both the use and effectiveness of OJS have changed dramatically since the turn of the twenty-first century. Kuhn and Mansour (2014) find that the unemployed in 2008–09 were three times more likely to use OJS than about ten years earlier and that using OJS significantly increased job seekers' chances of finding work relative to using traditional methods only. Using 2011 data, we confirm Kuhn and Mansour's (2014) findings regarding the prevalence and effectiveness of OJS.

In this article, we discuss the evolution of online job search and how it has become the principal method for finding employment and hiring workers. We then proceed to discuss the typical content of the OJS data and the advantages and disadvantages of using these data for research. We also compare the aggregate labor market patterns from OJS data with the aggregate patterns from published U.S. statistics. Through that comparison, we show that the OJS data generally capture the aggregate patterns of the U.S. labor market. Finally, we summarize some recent research that uses OJS data—such as studies examining questions regarding search effort and duration of search, the scope of search, search effort and unemployment insurance (UI) benefit extensions, and the incidence of wage posting online. We also briefly review recent studies that use the OJS market for experimental research on job search and hiring. By going over the recent literature, we hope to show how OJS data can provide exciting avenues for future research on the labor market.

The evolution of online job search

Using the Internet to find work is a fairly new phenomenon. Prior to the rise of OJS, job seekers had to rely on newspaper help-wanted ads, referrals, offline networking, word-of-mouth leads on jobs, and direct contact with employers to find out about new job opportunities. Kuhn and Skuterud (2004) were the first to look at how OJS compared with these traditional methods of job search. They examined the job-finding prospects of the unemployed using the 1998 and 2000 *Computer and Internet Use Supplements* to the *Current Population Survey* (CPS). Internet use in general, let alone for job search, was not yet common in 2000. As panel A of table 1 shows, less than 40 percent of the unemployed had Internet access at home back then. Consequently, only about one-quarter of the unemployed used the Internet for job search during that time (panel B of table 1). The fractions of the employed and those out of the labor force using the Internet for job search were even lower.

TABLE 1

Percentage having Internet access at home and performing online job search, by labor force status, 2000 and 2011

	A. Internet access at home		B. Performing online job search	
	2000	2011	2000	2011
Employed				
At work	52.1	82.7	11.3	38.1
Absent (on leave)	61.1	87.8	10.5	33.6
Unemployed				
On layoff	39.6	74.4	10.3	48.3
Looking for work	39.4	71.0	25.5	76.3
Not in the labor force				
Retired	23.8	59.3	0.5	7.8
Disabled	20.4	51.1	2.2	16.9
Other	45.7	76.7	6.3	32.6

Note: This table reports the percentage of each group having Internet access at home (panel A) or using the Internet to search for work online (panel B).
Sources: Estimates for 2000 from Kuhn and Skuterud (2004), based on data from the U.S. Bureau of Labor Statistics, 2000 *Computer and Internet Use Supplement* to the *Current Population Survey*; estimates for 2011 from authors' calculations based on data from the U.S. Bureau of Labor Statistics, 2011 *Computer and Internet Use Supplement* to the *Current Population Survey*.

TABLE 2

Probit estimates of job-finding probabilities within one year for the unemployed, 1998/2000 and 2011

A. 1998 and 2000 CPS *Computer and Internet Use Supplements*

Dependent variable	**1**	**2**	**3**	**4**
Internet job search	0.089	0.055	0.062	0.019
	(0.095)	(0.107)	(0.097)	(0.109)
Internet access at home		0.065		0.083
		(0.096)		(0.097)
2000 year dummy	–0.111	–0.118	–0.120	–0.129
	(0.073)	(0.074)	(0.074)	(0.075)
Search controls	No	No	Yes	Yes
Number of observations	1,344			

B. 2011 CPS *Computer and Internet Use Supplement*

Dependent variable	**1**	**2**	**3**	**4**
Internet job search	0.281**	0.260**	0.265**	0.244**
	(0.068)	(0.068)	(0.069)	(0.069)
Internet access at home		0.159*		0.175*
		(0.075)		(0.076)
Search controls	No	No	Yes	Yes
Number of observations	1,964			

* Significant at the 10 percent level
** Significant at the 5 percent level

Notes: This table reports the probit estimates of the probability of being employed one year later on dummies for whether an unemployed individual engaged in online job search or had Internet access at home. Panel A reports estimates using the 1998 and 2000 *Computer and Internet Use Supplements* to the *Current Population Survey* (CPS), while panel B reports estimates using the 2011 CPS *Computer and Internet Use Supplement*. Each column represents a different regression specification. All regressions are based on the probit models from Kuhn and Skuterud (2004) (see also note 1) and include controls for demographics (sex, four age categories, marital status, sex × marital status, five education categories, race/ethnicity, homeowner status, and immigrant status) and labor force status (unemployment duration, layoff status, four categories for previous employment or schooling status, and three categories for type of previous job—that is, private, public, or self-employed). Search controls include nine categories of the search method type(s) used. Standard errors are in parentheses.
Sources: Authors' calculations based on data from the U.S. Bureau of Labor Statistics, 1998, 2000, and 2011 *Computer and Internet Use Supplements* to the *Current Population Survey*.

Kuhn and Skuterud (2004) find that during 1998 and 2000, OJS did not increase the probability of finding work relative to traditional search methods. We repeat their analysis using the same 1998 and 2000 CPS *Computer and Internet Use Supplements*. We estimate a probit regression of the effect of OJS use by the unemployed on their probability of finding employment one year later. Panel A of table 2 shows that in 1998 and 2000, the use of Internet job search did not significantly increase the unemployed's chances of finding work a year later (see first row).[1]

Recent work by Kroft and Pope (2014) finds a similar pattern by analyzing online job postings on craigslist. These authors exploit the fact that craigslist entered different cities at different times—which allows them to explore the within-city variation in craigslist postings and the local unemployment rate over time. They find that the introduction of craigslist into a city had no effect on the local unemployment rate. This is despite the fact that the introduction of craigslist, which also lists ads for apartment rentals, among other things, significantly reduced local rental vacancy rates.

More-recent data, however, paint a very different picture regarding the effectiveness of OJS. Kuhn and Mansour (2014) use the CPS data along with data from the 1997 cohort of the *National Longitudinal Survey of Youth* (NLSY97) from the U.S. Bureau of Labor Statistics. They find that the use of OJS by the unemployed tripled between 1998/2000 (the years studied by Kuhn and Skuterud, 2004) and 2008–09. Furthermore, they find that using OJS in the later period reduced job seekers' unemployment duration by about 25 percent. Given the pervasiveness of OJS use in 2008–09, their results may reflect a penalty for those that rely solely on traditional search methods more than an increase in the effectiveness of OJS.

Our analysis confirms Kuhn and Mansour's (2014) findings. We apply the methodology of the original Kuhn and Skuterud (2004) study, but use more-recent data from the 2011 CPS *Computer and Internet Use Supplement*. Tables 1 and 2 contain our results. Panel B of table 1 (p. 2) shows that the use of OJS tripled among the unemployed (not on layoff), from 26 percent in 2000 to 76 percent in 2011. The use of OJS by the employed and those out of the labor force increased considerably as well, despite the fact that not all individuals in these groups were necessarily seeking work. Among the employed (either at work or on leave), the use of OJS rose from about 11 percent in 2000 to over one-third in 2011. Among those out of the labor force and not retired or disabled, OJS use increased from 6 percent in 2000 to nearly 33 percent in 2011. An important reason behind the prevalence of OJS use in 2011 was increased access to the Internet. Over 82 percent of the employed and at least 71 percent of the unemployed had Internet access at home in 2011 (see panel A of table 1).

More importantly, we find that in 2011, the use of OJS significantly increased the probability of finding a job within a year by an unemployed job seeker. Panel B of table 2 displays the results of our probit analysis based on the Kuhn and Skuterud (2004) study that we applied to 2011 data. The estimates in the first row of table 2, panel B show that Internet job search increases the unemployed's probability of landing a job within a year by about 25 percent—a finding in line with the results of Kuhn and Mansour (2014). These estimates run counter to the results using 1998/2000 data reported in panel A of table 2.

Online job search has certainly become a common and effective means of job search since 2000. In fact, one of the most widely used indicators of labor demand—the Conference Board Help Wanted OnLine (HWOL) Data Series—has its estimates generated by using only counts of online job postings.[2] Delving deeper into the ways in which OJS is conducted shows that job search websites, such as CareerBuilder and Monster, play a large role, but online job boards are not the only way in which the Internet has affected job search. Employers are now much more likely to require job seekers apply directly through company websites. Furthermore, the types of jobs one can apply for on the Internet represent a broad range of occupations. For instance, in our own research (Faberman and Kudlyak, 2016), we use data from an employment website that caters specifically to hourly jobs, which tend to be lower skilled and lower paying than salary jobs. There are a variety of employment websites that cater to specific occupations, industries, or types of jobs. In summary, OJS has become a mainstay within the U.S. labor market.

Online job search data: Typical content and representativeness of the labor market

As OJS has become an integral part of matching job seekers with potential employers, the data on OJS have provided new and exciting opportunities for research. In this section, we discuss the typical content of OJS data and go over the advantages and disadvantages of using these data for research. We then show that the data are generally representative of the broader labor market.

Recent research has used proprietary microdata obtained from OJS websites. These data contain a wealth of information on job seekers' search behavior and on the types of vacancies that employers post. The data provide opportunities for answering questions about job search and recruiting that other existing data sources, such as household surveys, cannot address.

One of the main advantages of OJS data is their size and scope. The microdata from a job posting website contain the universe of registered job seekers and the universe of vacancies posted on the website over a given period. Consequently, the data sets are often large enough to allow for very fine disaggregation by detailed geography, occupation, or other categorizations. In addition to their rich cross-sectional nature, OJS data feature a longitudinal dimension. Since OJS data contain information on actual (as opposed to self-reported) behavior, they provide detailed and accurate pictures of job seekers' timing of search, search effort, and scope of search, as well as the degree to which they target their search. The vacancy information typically includes the job title (which provides more details about the type of job than a standard occupational code[3]) and the job's requirements, including education and experience requirements. Additionally, vacancy postings sometimes feature wage details, at the employer's discretion (only a fraction of vacancies, generally between one-fifth and one-third, include information about the wage offered[4]). In the next section, we review research that studies the incidence of wage posting in online job ads. Finally, many online job boards require job seekers to set up online profiles, generating an additional source of information on the job seekers' backgrounds and employment histories.

To our knowledge, OJS data used in the existing studies do not contain information on the outcomes of job seekers' searches or employers' job vacancy postings. Often, such information is not available to the online job board because employers generally respond to an application outside of the online job board. In our research with OJS data (Faberman and Kudlyak, 2016), we use the first and last time a job seeker applies

FIGURE 1

Aggregate vacancies over time: Snagajob postings versus JOLTS data, September 2010–September 2011

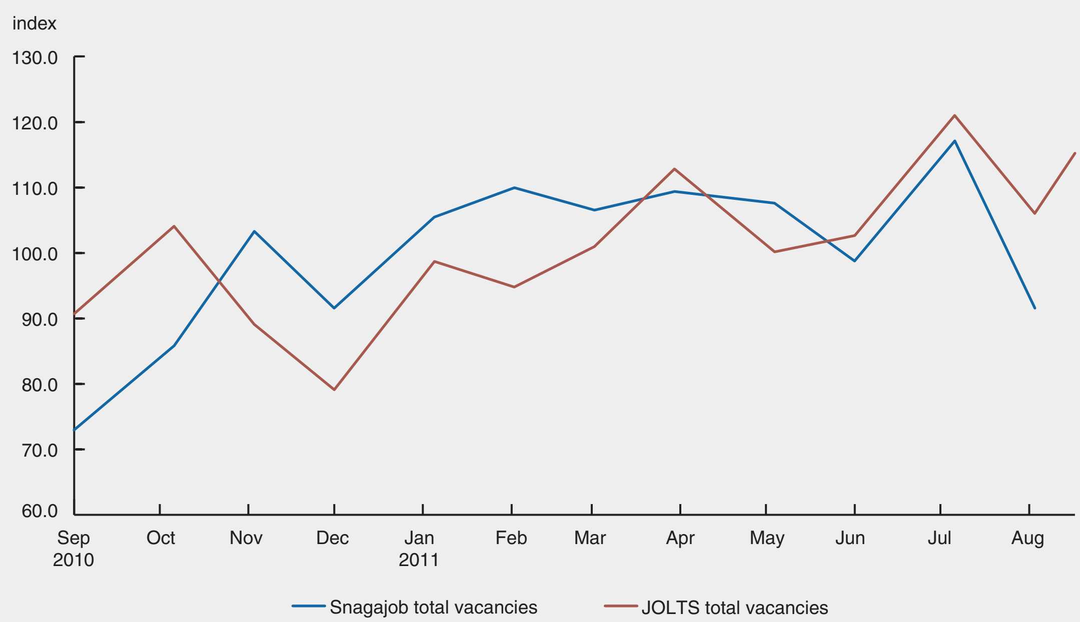

Notes: The JOLTS vacancies are nonfarm employment vacancies. The estimates for both series are *not* seasonally adjusted and are indexed such that the sample mean of each series equals 100.
Sources: Authors' calculations based on data from Snagajob; and U.S. Bureau of Labor Statistics, *Job Openings and Labor Turnover Survey* (JOLTS).

for a job on the website to infer the approximate length of an individual's job search and the first and last time we observe someone applying to a vacancy to infer the approximate duration of a job posting.

Because OJS data are limited to information on job search on the particular online job board, researchers do not have information on other methods or other websites that a job seeker is using to find work. Researchers also do not know whether the start of the search on the website coincides with the start of a broader job search or whether the end of the search on the website represents finding a job, switching to another search method, or ending all search efforts without finding work. That said, OJS data provide information on job search and hiring behavior that is unavailable in standard labor market data from surveys or administrative records.

A recurring and reassuring finding among researchers who work with OJS data is that the patterns observed in OJS data tend to match the aggregate patterns observed in published statistics of publicly available data, such as those from the federal government. Comparisons of time-series patterns cover a relatively short time span; however, comparisons of cross-sectional patterns are more straightforward and generally line up well with the estimates from publicly available data.[5]

We use data from the job search website Snagajob to show that its time-series patterns in 2010–11 are roughly consistent with the time-series patterns of published statistics.[6] Figure 1 compares the number of active vacancies in the Snagajob data with the number of vacancies reported in the *Job Openings and Labor Turnover Survey* (JOLTS) data from the U.S. Bureau of Labor Statistics. Both sets of data are indexed so that their respective time-series means equal 100. The figure shows that except for some differences at the beginning of the period, the two time series match up quite well. Similarly, figure 2 compares the

FIGURE 2

Aggregate job seekers over time: Snagajob nonemployed versus CPS unemployed, September 2010–September 2011

Notes: For the derivation of the nonemployed job seekers in the Snagajob data, see note 7. The estimates for both series are *not* seasonally adjusted and are indexed such that the sample mean of each series equals 100.
Sources: Authors' calculations based on data from Snagajob; and the U.S. Bureau of Labor Statistics, *Current Population Survey* (CPS).

TABLE 3

Demographic characteristics: Snagajob sample versus the CPS, September 2010–September 2011

	Share of Snagajob job seekers	Share of CPS unemployed	Share of the CPS labor force
Gender			
Male	43.2	56.3	53.3
Female	56.8	43.7	46.7
Age			
16–24 years old	48.5	26.3	13.6
25–39 years old	26.4	31.6	32.2
40–54 years old	18.1	27.4	34.2
55 years old and older	6.9	14.7	19.9
Education			
High school or less	58.4	51.0	37.1
Certification or some college	11.0	19.5	17.1
Associate's degree	14.3	20.0	10.6
Bachelor's degree or more	16.3	9.4	35.1
Race/ethnicity			
White	50.0	54.4	67.2
Black	26.1	19.4	11.0
Hispanic	14.2	19.2	14.8
Other	9.7	6.9	6.9

Notes: This table reports the percentage of individuals in each demographic category for a sample of online job seekers (both nonemployed and employed) searching on Snagajob for at least one week over the period September 2010 through September 2011 and for those reporting themselves as unemployed and the full labor force (both unemployed and employed persons) in the CPS (CPS statistics are monthly averages over the period September 2010 through September 2011). Some columns may not total because of rounding.
Sources: Authors' calculations based on data from Snagajob; and the U.S. Bureau of Labor Statistics, *Current Population Survey* (CPS).

number of nonemployed job seekers in the Snagajob data with the number of unemployed individuals in the CPS.[7] Again, except for some differences in the first months of the sample period (likely due to the difficulty in identifying the nonemployed in the Snagajob data at the beginning of the period), the two series line up quite well.

We also use the Snagajob data to compare the demographics of online job seekers with the demographics of the unemployed and the full labor force in the CPS. The Snagajob website focuses on hourly wage jobs, which tend to be lower paying than salary jobs, typically require less skill and education than salary jobs, and tend to be filled by young workers. Consequently, the job seekers using the website are not expected to be representative of the U.S. labor force as a whole. We focus on individuals (both nonemployed and employed) who search on the website for at least one week. Our estimates are in table 3. Table 3 shows that the Snagajob job seekers are not too different from the unemployed in the CPS in terms of their demographics as of 2010–11. This is partly because the unemployed tend to be younger and lower skilled relative to the labor force as a whole. Comparing the Snagajob job seekers to the unemployed in the CPS shows that the former tend to be younger and are more likely to be female, but are otherwise similar (in terms of their race/ethnicity and educational attainment).

Finally, we use the Snagajob data to examine how job search spells of the nonemployed on the website (see note 7) compare with the duration of unemployment calculated from the CPS. Table 4 compares

TABLE 4

Differences in job search and unemployment duration: Snagajob sample versus the CPS, July 2011

	Snagajob nonemployed job seekers	CPS unemployed
Duration (percent)		
Less than five weeks	52.5	20.5
Five–14 weeks	38.0	24.2
15–26 weeks	7.0	12.2
27 or more weeks	2.5	43.1
Mean duration (weeks)	6.3	39.0
Median duration (weeks)	4.0	19.7

Notes: The first four rows of this table report the share of nonemployed online job seekers on Snagajob (see note 7) who submitted more than one application, as well as the unemployed in the CPS, with an active search spell within the listed range. The summary statistics on the duration of (incomplete) search and unemployment spells are in the final two rows. The estimates for nonemployed online job seekers are based on a cross section of Snagajob job seekers identified as actively searching during the CPS reference week of July 2011. The CPS reference week is usually the calendar week (Sunday through Saturday) containing the 12th day of the month.

Sources: Authors' calculations based on data from Snagajob; and the U.S. Bureau of Labor Statistics, *Current Population Survey* (CPS).

the distribution of search spells of nonemployed job seekers (with more than one job application) from the Snagajob data to the distribution of unemployment spells in the CPS. Both sets of estimates are for a cross section of job seekers in July 2011, implying that the estimates are for incomplete search spells and unemployment spells (that is, spells in progress). First, table 4 shows that the nonemployed's search spells online according to the Snagajob data are considerably shorter than unemployment spells according to the CPS. This is not too surprising because search on a particular website is generally one of several methods the unemployed use to look for work. The median search spell of the nonemployed on Snagajob is about four weeks, while the median unemployment spell in the CPS is just under 20 weeks. About 53 percent of online search spells of the nonemployed on Snagajob last one month or less, while about 21 percent of unemployment spells reported in the CPS last one month or less. Second, the biggest difference between the distribution of search spells of the nonemployed on Snagajob and the distribution of unemployment durations in the CPS is in the fraction of long-duration job seekers (that is, those searching for work more than six months). In July 2011, over 43 percent of the CPS unemployed had been searching for that long, but only about 3 percent of nonemployed job seekers on Snagajob had been searching that long.[8]

Recent studies of job search and hiring

We now turn to a brief overview of recent research conducted using OJS microdata. The studies are chosen to showcase the range of questions studied using OJS data.

Brown and Matsa (2012) use OJS microdata from the CareerBuilder website to examine the relationship between an employer's financial health and its recruiting outcomes. They focus on over 5.5 million job applications to 40 financial firms during the Great Recession period (2008–10). Brown and Matsa estimate firms' financial health through credit default swap (CDS)[9] prices. A higher CDS price implies a higher probability that the firm may default on its debt and, therefore, signals weaker financial health. Between 2008 and 2010, many firms saw large swings in their respective CDS prices, providing the authors with the variation needed to identify any effects. Using survey data, the authors first show that job seekers can correctly infer when a firm is in financial distress. They then show that a 10 percentage point increase in a firm's probability of default (as measured by its CDS price) results in about a 20 percent decline in the number of applications. This effect is driven predominantly by changes in the labor supply; that is, after inferring the financial health of a distressed firm, job seekers choose to send it fewer applications because of its higher default risk. It is not the case that the composition of job openings changes or that the wage changes in response to the change in default risk. Finally, Brown and Matsa show that those who do apply to a financially distressed firm tend to be applicants from the zip codes with lower average earnings and education.

Brenčič (2012) is among the first to use OJS data to study the incidence of wage posting.[10] The data come from the employment website Monster in the United States. She also examines traditional job postings by using data from government employment agencies in the United Kingdom and Slovenia. Across all data sources, she finds that employers are less likely to post a wage offer when searching for highly skilled workers. In particular, Brenčič finds that employers are more likely to post a wage offer in their job ads when they are less concerned about drawing an adversely affected pool of job applicants—that is, when 1) a recruitment agency screens the job applicants (reducing the chances of the employer making a bad hire); 2) the opportunity costs of search are high (for example, when the job needs to be filled immediately) and thus employers are less selective; and 3) employers aim to hire workers with fewer skills and these skills are easier to measure.

In our own research (Faberman and Kudlyak, 2016), we use data from Snagajob to examine the relationship between search effort and search duration. We examine the search behavior of all job seekers on the employment website over the course of about one year (from September 2010 through September 2011). The sample period provides us with over 17 million applications from roughly 4.8 million job seekers for 1.4 million vacancies. We estimate a job seeker's search effort as the number of applications sent out each week. Since the data set is large and longitudinal, we are able to control for a variety of observable and unobservable characteristics.

In Faberman and Kudlyak (2016), we focus on controlling for two main factors that can affect the relationship between search effort and search duration: 1) the composition of job seekers and 2) the number of existing and newly posted vacancies (that is, the stock and flow of vacancies). In theory, job seeker composition could matter in determining this relationship because individuals who exert fairly low effort throughout their searches are less likely to find a job. Several studies of the unemployed (for example, Krueger and Mueller, 2011) find that search effort declines with duration. If individuals in a given sample differ in the amount of effort they put forth over the duration of their searches, then this declining relationship could be explained in part by the job seeker composition of the sample; that is, the long-duration job seekers are also the low-effort job seekers, which would explain why search effort declines for the sample, and it would not necessarily be the case that all individuals' search effort falls over time. How job seekers search through the relevant range of vacancies over time can also affect whether search effort declines with duration. For example, if their search pattern is "stock-flow" in nature, as in the model of Coles and Smith (1998), at the beginning of their searches job seekers search through the entire stock of vacancies, but afterward they examine only new vacancies that come along (that is, the flow of new vacancies). This type of search behavior can cause a steep decline in the number of applications that job seekers send after they search through the initial stock of vacancies.

Consistent with the existing studies, we find that individual search effort declines with duration, after controlling for a variety of factors, in our recent research. We find that controlling for the (unobserved) heterogeneity among job seekers is important for this result, although controls for a stock-flow nature of search do not alter our conclusions. Moreover, we find that long-duration job seekers typically send more applications in every period of their search than short-duration job seekers—which runs counter to standard search models that include job seeker search effort. Figure 3 replicates this finding from Faberman and

FIGURE 3

Applications per week over the duration of job search, by completed search spell length

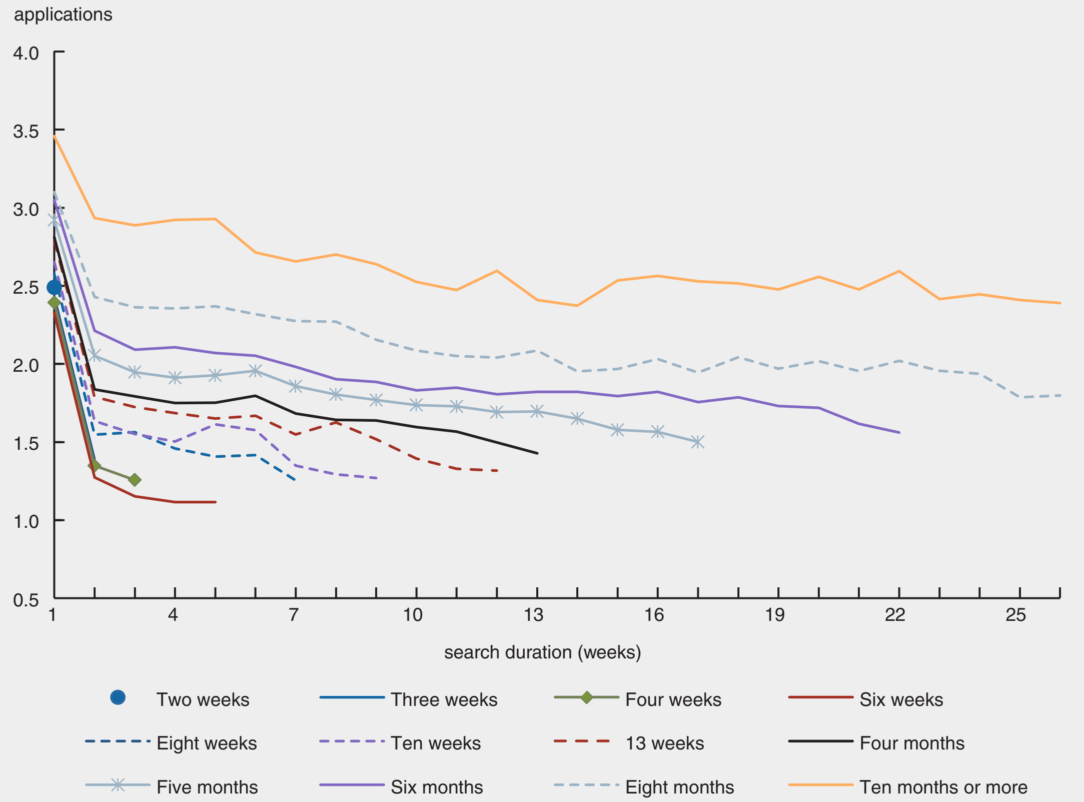

Notes: The estimates here are from our recently revised working paper (see the source line) using data from September 2010 through September 2011. This figure shows the estimated relationship between applications per week and the duration of search separately for job seekers based on the total length of their completed search spells. The estimates control for observed demographics of the job seeker, fixed effects for the starting month of the search spell and the job seeker's metropolitan area, and the number of new and existing (metro area) vacancies active during the week of search. By design, the last week of the search spell is excluded from the figure (see our working paper for further details). Only selected search spell lengths are reported.
Source: Faberman and Kudlyak (2016).

Kudlyak (2016). The figure shows search effort (measured by applications sent each week) as a function of the number of weeks since the beginning of the search while controlling for a variety of observable characteristics of the job seekers and their respective local labor markets. The figure displays the relationship separately by duration of completed search spells. Relative to all other groups in our sample, the job seekers whose completed search spells lasted ten months or more typically send more applications per week throughout their search spells. As the search spell continues, the average number of applications per week falls for all job seekers, but the average number of applications per week increases monotonically with the duration of completed search spell. Those who search for three months or less, for example, send considerably fewer applications per week throughout their search spells than those who search for more than three months.

In Faberman and Kudlyak (2016), we then show that the high observed effort of the long-duration job seekers using Snagajob is consistent with the notion that those with the lowest expected returns to search exert the highest effort. It is analogous to a dominant income effect in the labor supply decision (that is, workers opt to supply more labor when a wage decline makes them feel poorer). In particular, we find that search effort is highest in metropolitan areas with weak labor markets, as well as among job seekers who tend to be nonemployed and older and tend to have separated from a long-tenured (that is, stable) job. These job seekers seem to be aware of their poor prospects and attempt to compensate for them with greater search effort than others.

In a series of papers, Ioana Marinescu and her co-authors use data from CareerBuilder to study the behavior of job seekers and firms. Marinescu (2014) uses these data from 2007 through 2011 to examine the effect of unemployment insurance benefit extensions on job-seeker search effort. Using the state-level time-series variation in the incidence of UI benefit extensions during the Great Recession, she shows that a one-week extension of UI benefits leads to a 0.4 percent reduction in the number of applications sent, and has essentially no effect on the number of vacancies posted. Marinescu and Rathelot (2014) use the CareerBuilder data to study geographic mismatch—that is, the misallocation of jobs and job seekers across space—in 2012. They find that geographic mismatch can only explain about 5 percent of unemployment (which we calculate to be equivalent to about 0.3 percentage points of the 2012 unemployment rate). This is mainly because job seekers send only a small fraction of their applications outside their local labor market and because they are already located where the bulk of the vacancies are. Finally, Marinescu and Wolthoff (2015) use the job titles contained in the posted vacancy descriptions in the CareerBuilder data to examine whether these titles convey any information about the position to potential applicants. They find that the job titles explain about 90 percent of the variation in posted wages. They also find, however, that only a small fraction of vacancies (20 percent) actually post any wage data. Job titles also account for about 80 percent of the variation in the education and experience levels of the applicants for a job vacancy. One interpretation of these findings is that the job title conveys nearly as much information about the job as a posted wage would.

Similar to Marinescu (2014), Baker and Fradkin (2015) use OJS search data to examine the effects of unemployment insurance benefits on job search effort. These researchers use data from Google searches (containing the term "jobs") and find that individuals using the search engine to find work while receiving UI benefits search 30 percent less than those doing the same while not on UI. Furthermore, those who are about to exhaust their UI benefits search twice as much as claimants with more than 30 weeks of such benefits left. Despite these differences in search effort, they find that the extension of UI benefits during the Great Recession had only a modest effect on the unemployment rate.

Kudlyak, Lkhagvasuren, and Sysuyev (2014) use the Snagajob microdata to examine how the scope of a job seeker's search evolves over time. They find that as the search progresses, job seekers tend to apply to lower-quality jobs, as measured using the educational attainment of the other applicants to the same job. The authors argue that the evidence is consistent with a declining reservation wage (the lowest wage rate at which a job seeker would accept employment) over time.

Banfi and Villena-Roldán (2016) use microdata for Chile from the Latin American OJS website Trabajando.com to test for evidence of directed search—as in the model of Shimer (1996) or Moen (1997). In models of directed search, firms that post vacancies with a higher wage than other job postings attract more applicants and fill those vacancies more quickly, all else being equal. Banfi and Villena-Roldán's data contain information on the wages that firms expect to pay regardless of whether they post them in the job ads. The authors find evidence consistent with directed search, even when a wage is not explicitly mentioned in the job posting. They conclude that the other information included in the job posting (such as job requirements) conveys something about the expected wage, consistent with the evidence of Marinescu and Wolthoff (2015). Kuhn and Shen (2013) use OJS data from China, where it is legal to list an explicit gender preference on a job posting, to test for gender discrimination. They find that women are preferred as often as men in China's gender-targeted online job ads, but usually on the basis of beauty, youth, and height rather than skills. Helleseter, Kuhn, and Shen (2016) follow up with a study of both gender and age discrimination using OJS data from China and Mexico. They find that gender preferences "twist" sharply away from women toward men when a preference for an older worker (over a young one) is listed.

Two studies—Modestino, Shoag, and Ballance (2015) and Hershbein and Kahn (2015)—use data on online vacancy postings from Burning Glass Technologies, a company that collects data on vacancies posted across an array of OJS websites. They examine whether the skill requirements of the vacancies posted change in response to shocks to the labor market. Modestino, Shoag, and Ballance (2015) examine whether a positive shock to the local labor supply—in their case, driven by the withdrawals of troops from Iraq and Afghanistan—leads to "upskilling" (that is, an increase in skill requirements) in jobs most likely to be sought by returning veterans. They find that up to 30 percent of the observed upskilling in their data can be attributed to these labor supply shocks during the period 2007–10. Hershbein and Kahn (2015) examine whether slack labor markets lead to higher skill requirements between 2010 and 2014, and they find that increases in the local unemployment rate lead to increases in the levels of education and experience required for posted vacancies, particularly for occupations with routine tasks.

Finally, several studies have used OJS websites for experimental research. In an influential study, Kroft, Lange, and Notowidigdo (2013) examine whether there is a stigma associated with long-term unemployment by submitting resumes to potential employers in response to job postings online. Resumes are made to be identical save for the amount of time listed since the last job. The authors find that callback rates are lower when unemployment duration is longer, and take their findings as evidence that employers use unemployment duration as a criterion for screening job applicants.

Pallais and Sands (2016) use an online job board to actually hire individuals for short-term independent contractor work. They examine whether referrals for these individuals affect their performance outcomes. The researchers find that the referrals are a positive signal of the productivity of these contractors (relative to contractors without referrals), independent of their on-the-job interactions with their referrers. They conclude that referrals convey information about worker quality that is not contained in other signals an employer may receive.

Gee (2014) examines application behavior on LinkedIn, the professional networking website, to see if having additional information on the number of people who started to apply for a particular job affects a job seeker's probability of applying herself. After randomizing which applicants see the number of applicants to a position, Gee finds that the additional information does in fact increase the chances a job seeker starts or finishes the application.

Finally, Belot, Kircher, and Muller (2015) develop a web interface for an OJS website that provides job search advice to a random subsample of job seekers. The advice is tailored to the job seekers' labor market characteristics. They find that advice from the new web interface leads job seekers in the subsample to search through a broader range of jobs and to experience a higher callback rate (for interviews) relative to those who were not provided this advice. These results are driven mainly by those who were unemployed for several months and those who were initially conducting a narrow search.

Conclusion

Online job search has become an integral part of the labor market. In fact, the use of online job search websites by job seekers and employers has essentially become the norm. A comparison of early research on OJS with more-current studies shows that its pervasiveness and effectiveness among job seekers have increased substantially since 2000. Today, job seekers who use online job search as one of their search methods are much more likely to find work—and find work faster—than those who do not.

Furthermore, the OJS websites and the data available from them have opened up new avenues for research on the labor market. In many instances, the data sets are large, longitudinal, and generally representative of the broader labor market. Recent research using OJS data has focused on topics such as the search effort of job seekers, the hiring behavior of firms, the role of posted wages and job descriptions in attracting (different types of) workers, and the role of aggregate labor market conditions in determining job requirements. Current research has also used OJS websites as settings for experimental research on the labor market. These have included studies that examine the stigma effects of unemployment duration, the role of referrals in hiring, and the role of additional information in targeting or improving job search. In short, online job search has become the primary channel by which workers and employers are matched, and data on online job search have the potential to help understand this matching process better than ever before.

NOTES

[1]Note that, on account of minor differences in our data samples and regression specifications, we obtain slightly different estimates than the original Kuhn and Skuterud (2004) study.

[2]For background information on HWOL (and the release analyzing its latest results), see https://www.conference-board.org/data/helpwantedonline.cfm. The Conference Board decided to stop publishing HWOL's predecessor, the Help-Wanted Advertising Index (of print advertising) in July 2008, largely because the organization determined this index no longer adequately captured labor market demand; for details, see Conference Board (2008).

[3]See Marinescu and Wolthoff (2015) for more information on the job titles of vacancies posted online.

[4]See, for example, Brenčič (2012) and Marinescu and Wolthoff (2015). In many cases, the relative lack of posted wage information in job postings may be for strategic reasons (see our discussion of Brenčič, 2012, in our literature review section), or may be due to the fact that the wage is negotiable, as Hall and Krueger (2012) find.

[5]For example, see Marinescu and Wolthoff (2015), who show that their data on posted wages from CareerBuilder align well with the wage distribution derived from the *Current Population Survey*.

[6]See Faberman and Kudlyak (2016) for more details on the Snagajob data.

[7]We identify the nonemployed (that is, all individuals without a job) in the Snagajob data by using the employment history records that they report on the website. For those without employment history records on the site, we simply extrapolate on them the monthly fraction of the nonemployed among Snagajob job seekers with such records. (For details on how the federal government makes distinctions among the nonemployed for the CPS—for example, how it separates them into the unemployed and those out of the labor force—see http://www.bls.gov/cps/cps_htgm.htm. Such distinctions do not apply to the Snagajob data.)

[8]Note that it is possible, however, that the long-term unemployed reported in the CPS are represented by a sequence of multiple short-term search spells in the online job search data from Snagajob. See Faberman and Kudlyak (2016) for more details.

[9]For a definition of CDS, see http://lexicon.ft.com/Term?term=credit-default-swap--CDS.

[10]See Hall and Krueger (2012) for a study about the prevalence of wage posting versus wage bargaining, using survey data.

REFERENCES

Baker, Scott R., and Andrey Fradkin, 2015, "The impact of unemployment insurance on job search: Evidence from Google search data," Northwestern University, mimeo, November 7, http://andreyfradkin.com/assets/FullTexasJobSearch.pdf.

Banfi, Stefano, and Benjamín Villena-Roldán, 2016, "Do high-wage jobs attract more applicants? Directed search evidence from the online labor market," University of Chile, Department of Industrial Engineering, Center for Applied Economics, mimeo, April 11, http://www.benjaminvillena.com/data/uploads/Directed%20Search%20evidence.pdf.

Belot, Michèle, Philipp Kircher, and Paul Muller, 2015, "Providing advice to job seekers at low cost: An experimental study on on-line advice," University of Edinburgh, School of Economics, discussion paper, No. 262, November, http://www.econ.ed.ac.uk/papers/id262_esedps.pdf.

Brenčič, Vera, 2012, "Wage posting: Evidence from job ads," *Canadian Journal of Economics*, Vol. 45, No. 4, November, pp. 1529–1559.

Brown, Jennifer, and David A. Matsa, 2012, "Boarding a sinking ship? An investigation of job applications to distressed firms," National Bureau of Economic Research, working paper, No. 18208, July, http://www.nber.org/papers/w18208.

Coles, Melvyn G., and Eric Smith, 1998, "Marketplaces and matching," *International Economic Review*, Vol. 39, No. 1, February, pp. 239–254.

Conference Board, 2008, "The Conference Board to discontinue help-wanted print advertising index," PR Newswire, April 4, http://www.prnewswire.com/news-releases/the-conference-board-to-discontinue-help-wanted-print-advertising-index-57104497.html.

Faberman, R. Jason, and Marianna Kudlyak, 2016, "The intensity of job search and search duration," Federal Reserve Bank of Richmond, working paper, No. 14-12R, revised February 2016, https://www.richmondfed.org/publications/research/working_papers/2014/wp_14-12r.

Gee, Laura K., 2014, "The more you know: Information effects in job application rates by gender in a large field experiment," Tufts University, Department of Economics, working paper, September 17, http://ase.tufts.edu/economics/documents/papers/2014/geeMoreYouKnowNOappendix.pdf.

Hall, Robert E., and Alan B. Krueger, 2012, "Evidence on the incidence of wage posting, wage bargaining, and on-the-job search," *American Economic Journal: Macroeconomics*, Vol. 4, No. 4, October, pp. 56–67.

Helleseter, Miguel Delgado, Peter Kuhn, and Kailing Shen, 2016, "Age and gender profiling in the Chinese and Mexican labor markets: Evidence from four job boards," National Bureau of Economic Research, working paper, No. 22187, April, http://www.nber.org/papers/w22187.

Hershbein, Brad, and Lisa B. Kahn, 2015, "Is college the new high school? Evidence from vacancy postings," Yale University, mimeo, March 30, http://faculty.som.yale.edu/lisakahn/documents/HershbeinKahn_3_30_2015_draft.pdf.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo, 2013, "Duration dependence and labor market conditions: Evidence from a field experiment," *Quarterly Journal of Economics*, Vol. 128, No. 3, pp. 1123–1167.

Kroft, Kory, and Devin G. Pope, 2014, "Does online search crowd out traditional search and improve matching efficiency? Evidence from Craigslist," *Journal of Labor Economics*, No. 32, No. 2, April, pp. 259–303.

Krueger, Alan B., and Andreas Mueller, 2011, "Job search, emotional well-being, and job finding in a period of mass unemployment: Evidence from high-frequency longitudinal data," *Brookings Papers on Economic Activity*, Vol. 42, No.1, Spring, pp. 1–81.

Kudlyak, Marianna, Damba Lkhagvasuren, and Roman Sysuyev, 2014, "Systematic job search: New evidence from individual job application data," Federal Reserve Bank of Richmond, working paper, No. 12-03R, revised September 2014, https://www.richmondfed.org/publications/research/working_papers/2012/wp_12-03r.

Kuhn, Peter, and Hani Mansour, 2014, "Is Internet job search still ineffective?," *Economic Journal*, Vol. 124, No. 581, December, pp. 1213–1233.

Kuhn, Peter, and Kailing Shen, 2013, “Gender discrimination in job ads: Evidence from China,” *Quarterly Journal of Economics*, Vol. 128, No. 1, February, pp. 287–336.

Kuhn, Peter, and Mikal Skuterud, 2004, “Internet job search and unemployment durations,” *American Economic Review*, Vol. 94, No. 1, March, pp. 218–232.

Marinescu, Ioana, 2014, “The general equilibrium impacts of unemployment insurance: Evidence from a large online job board,” University of Chicago, mimeo, http://www.marinescu.eu/Marinescu_UI_2014.pdf.

Marinescu, Ioana, and Roland Rathelot, 2014, “Mismatch unemployment and the geography of job search,” University of Warwick, mimeo, December 1, http://www.sole-jole.org/15260.pdf.

Marinescu, Ioana, and Ronald Wolthoff, 2015, “Opening the black box of the matching function: The power of words,” Institute for the Study of Labor (IZA), discussion paper, No. 9071, May, http://ftp.iza.org/dp9071.pdf.

Modestino, Alicia Sasser, Daniel Shoag, and Joshua Ballance, 2015, “Upskilling: Do employers demand greater skill when workers are plentiful?,” Northeastern University, mimeo, January 21, http://www.northeastern.edu/cssh/wp-content/uploads/2014/08/Modestino-Shoag-and-Ballance_012114.pdf.

Moen, Espen R., 1997, “Competitive search equilibrium,” *Journal of Political Economy*, Vol. 105, No. 2, April, pp. 385–411.

Pallais, Amanda, and Emily Glassberg Sands, 2016, “Why the referential treatment? Evidence from field experiments on referrals,” *Journal of Political Economy*, forthcoming.

Shimer, Robert, 1996, “Contracts in a frictional labor market,” Massachusetts Institute of Technology, mimeo.

*R. Jason Faberman is a senior economist in the Economic Research Department at the Federal Reserve Bank of Chicago. Marianna Kudlyak is a senior economist in the Economic Research Department at the Federal Reserve Bank of San Francisco.**The authors thank Jacob Berman for excellent research assistance.*

ISSN 0164-0682