Working Paper Series

The Intensity of Job Search and Search Duration

WP 14-12

R. Jason Faberman
Federal Reserve Bank of Chicago

Marianna Kudlyak
Federal Reserve Bank of Richmond

THE FEDERAL RESERVE BANK OF RICHMOND
RICHMOND ▪ BALTIMORE ▪ CHARLOTTE

The Intensity of Job Search and Search Duration

May 2014

R. Jason Faberman, Federal Reserve Bank of Chicago
Marianna Kudlyak, Federal Reserve Bank of Richmond

Working Paper No. 14-12

Abstract

We use micro data on applications to job openings by individuals on a job search website to study the relationship between search intensity and search duration. Our data allow us to control for several factors that can affect the measured relationship between intensity and duration, including the composition of job seekers and changes in the number of available job openings over the duration of search. We find that a job seeker sends fewer applications per week as search continues. We also find that job seekers who search on the website longer tend to send more applications in every period. We attribute this finding to job seeker heterogeneity. Controlling for the local stock of vacancies does little to affect the result, mainly because job seekers continue to apply to older vacancies well into their search spell.

Keywords: Job applications, vacancies, labor market search effort, search duration.
JEL Codes: E24, J31, J24.

Faberman: jfaberman@frbchi.org, Kudlyak: mkudlyak@gmail.com. *We thank Fabian Lange, Melinda Pitts, Andreas Mueller, Ija Trapeznikova, Paul Beaudry, Giuseppe Moscarini, seminar participants at Notre Dame, the Kansas City Fed, and the Chicago Fed, and participants at several conferences for helpful comments. The authors are thankful to Ann Macheras at the Federal Reserve Bank of Richmond and to Scott Hicks and Amanda Richardson at SnagAJob for providing access to the data and gratefully acknowledge their comments and suggestions. The views expressed here are of the authors and do not reflect the views of the Federal Reserve Bank of Chicago, the Federal Reserve Bank of Richmond, or the Federal Reserve System.*

1. Introduction

A large literature documents a negative relationship between unemployment duration and the probability of finding a job (Jones, 1988; Machin and Manning, 1999; van den Berg and van Ours, 1996). Several factors have been put forth as causes of this negative duration dependence. Economists have pointed to the depreciation of human capital over the course of job search and to ex ante job seeker heterogeneity as potential causes of the declining probability of finding a job over the duration of search .[1] Variations in the search intensity of job seekers over the duration of a job seeker's search can also be an important factor for explaining negative duration dependence in job finding probabilities. While there are models that incorporate a role for search intensity, little empirical evidence exists on the relationship between search intensity and search duration.

In this paper, we use a large, novel set of longitudinal micro data on the application behavior of job seekers from an online job search engine to examine the relationship between search intensity and search duration. Our data include all applications made by job seekers to vacancies on the website at the daily frequency over one year. The data are nationally representative and include basic information on the geography and characteristics of both job seekers and vacancies. Our findings are twofold. First, using the weekly number of applications sent on the engine as our measure of search intensity, we find that an individual job seeker sends fewer applications per week as search continues. Second, we find that job seekers who search on the website longer, on average, send *more* applications per week throughout the length of their search spell than those who end their search on the website earlier.

Our results arise after addressing several measurement issues that our data are uniquely suited to deal with. One issue is that unobserved characteristics can affect the probability of finding a job and

[1] The studies include models of unemployment duration as a signal of worker productivity (Blanchard and Diamond, 1994), models of unemployment duration as a stigma (Lockwood, 1991; and Pissarides, 1992; with empirical support by Addison and Portugal 1989; and Kroft, Lange, and Notowidigdo, 2012), and models of heterogeneity in the employability of job seekers (Hornstein, 2012).

generate a relationship between search intensity and duration through a selection effect. For example, job seekers who search most intensely may find work faster. The longitudinal panel nature of our data allows us to control for unobserved job seeker characteristics through the use of fixed effects. Prior to these controls, we obtain a relationship between search effort and duration that is declining at first and slightly increasing thereafter. When including these controls, we find that search effort declines monotonically with duration. Another issue is that search intensity may vary with the stock of available vacancies. For example, Coles and Smith (1998) model job search as a stock-flow process. That is, at the start of their search job seekers search through the existing stock of vacancies but, thereafter, only search through newly-posted vacancies. Our data contain information on all vacancies ever applied to on the website. This allows us to construct measures of the stock of active vacancies within a job seeker's labor market. It also allows us to distinguish between newly-posted and pre-existing vacancies. Adding these controls, however, does little to affect our key findings. This is because most applications are to pre-existing vacancies throughout a job seeker's search spell. The fraction of applications to a newly-posted vacancy rises with duration, consistent with a stock-flow model, but it does so only slightly, representing only 17 percent of applications during a job seeker's sixth month of search. In comparison, job seekers send a similar fraction of their applications in the sixth month to vacancies that they applied to previously.

The data include rich detail on the location of each job seeker and each job opening, as well as basic demographic and industry information for each, respectively. The online data, however, do not have information on the outcome of any job application. Consequently, we do not know whether a job seeker stopped searching because they found work, they switched to another search method, or they stopped searching altogether. To address this issue, we replicate our analysis using a sample of "potential matches" that consists of job seekers who applied to an expiring vacancy the same week we identify as the end of their search spell. We provide compelling evidence that this sample reflects individuals who

likely found jobs through the website. Our two key results hold, with estimates that are very similar to those we obtain with our full sample. Other tests of robustness reject the notion that our findings are an artifact of narrowly examining search on a single website. Through a simulation where job seekers only differ in their tastes for search on the website, we show that job seeker attrition due to these tastes alone cannot account for our results. Furthermore, we find that the number of applications sent by a job seeker per week declines with duration in second and subsequent search spells among job seekers whom we identify as having multiple search spells on the website. This rejects the notion that the number of applications declines with duration as a job seeker learns how to better navigate the website and become more efficient in their search.

We also find that the number of applications sent by a job seeker per week is significantly higher in metropolitan areas with more slack labor markets. The finding is consistent with the idea that job seekers exert more effort throughout their search spell when they perceive their chances of finding employment as relatively weak. Finally, we examine the amount of time spent on the website per week and the applications sent per webpage viewed (a measure of application selectivity). The evidence suggests that longer-duration job seekers are less "choosy" in their job search, and that there is a strong correlation between applications per week and weekly time spent on the website, so that applications are likely a suitable measure of search effort.

Our finding of a decline in search intensity with duration rejects a simple search and matching model with a constant search effort. The finding suggests a need for richer search models that generate job seeker behavior that varies with the duration of search. These models might, for example, include changes in a job seeker's outside option, job seekers learning about their own skills (or the depreciation of their skills), or job seekers learning about the state of the labor market.

Our finding that long-duration job seekers exert the most search effort, at first look might appear at odds with search theoretic models that imply a positive effect of search effort on the probability of finding a job. At closer examination, it underscores the importance of job seeker heterogeneity. In particular, we interpret the finding as evidence that unobserved heterogeneity in job seeker characteristics is quantitatively more important than differences in search effort in explaining why some job seekers have long search spells. The heterogeneity may take a form of heterogeneity in employability or heterogeneity in the job seeker's search efficiency. First, long-duration job seekers may be viewed by potential employers as being low-productivity or otherwise unqualified for the job. These job seekers may exert greater search effort to increase their chances of finding a job. Second, long-duration job seekers may consistently apply to jobs that they are poorly suited for, so that, despite their substantially high effort, their job-finding probability remains low. This can occur due to imperfect information about a job seeker's own ability or about the labor market.

Our paper contributes to the literature that studies the relationship between job seeker search behavior and job-finding outcomes. The pioneering study by Krueger and Mueller (2011), which focuses on search effort and unemployment duration, is the most related to our own. In contrast to ours, they examine search effort using survey data on a cohort of unemployed workers in New Jersey. They conclude that the unemployed spend less time searching for work as search continues. It is not clear, however, how much measurement issues related to the repeated use of their survey instrument for reported time spent on job search affects their conclusions. More broadly, our work is related to studies that examine the relationship between the job-finding rate unemployment duration. The salient finding among these studies is the presence of negative duration dependence for the exit rate from unemployment. Finally, our paper is related to the growing literature that studies online job search (Kuhn and Skuterud, 2004; Kroft and Pope, 2012; Brown and Matsa, 2012; Marinescu and Wolthoff,

2012), and to the literature that studies hiring and vacancies (Barron, Bishop, and Dunkelberg, 1985; Holzer, Katz, and Krueger, 1991; Davis, Faberman, and Haltiwanger, 2013).

The rest of the paper is structured as follows. Section 2 describes the data, measurement, identification and provides basic evidence. Section 3 presents our main results on application behavior by duration. Section 4 presents results for the potential matched sample. Section 5 investigates the role of local demand. Section 6 presents robustness exercises and Section 7 concludes.

2. Data and Measurement

2.A. Data and Sample

We use proprietary data from SnagAJob, an online private job search engine. The data contain information about job seekers and their applications to online job postings (vacancies). Job seekers can browse information on the website at no cost. To apply for a job, however, a job seeker must register with the website. Registration is free. At registration, job seekers provide their basic demographic and geographic location information. Once registered, job seekers are able apply for posted job openings on the website by clicking an application link associated with the online job posting (which we refer to as a vacancy).

Employers pay a fee to post vacancies on the website on a per-period rather than a per-vacancy basis.[2] Each vacancy posting contains basic information on the geographic location and sector of the job. Postings also contain a general description of job duties that can vary widely in its detail on skills required, compensation offered, etc., though the dataset that we obtain from the website does not

[2] First time customers to the website may also post a single vacancy for free for up to 10 days. In addition to the limited duration, the free trial period does not provide for some of the other services, such as the opportunity to screen applicants with an assessment test.

contain this information. Most job postings provide at least basic information on job duties (including a job title, or occupation), skills required, and hours, including whether the job is full time or part time. Listed skill requirements can vary from detailed education, experience, and certification requirements, to only requiring that one be "highly motivated" or "committed to quality customer service."

An important feature of the job search engine is that the jobs posted there are predominantly hourly paid jobs. Hourly jobs tend to be lower-skilled jobs and tend to be concentrated in services. These jobs have been growing in importance and constitute a major part of the U.S. labor market (Autor et al., 2003; Acemoglu and Autor, 2012). Thus, the data provide a unique opportunity because, while the job seekers on the website tend to be lower-skilled and less educated than the general population, they are more representative of the unemployed population.

Our dataset contains daily micro data on all application-vacancy matches between September 2010 and September 2011, including detailed information on both applicants and vacancies during this period. We also have supplemental information on application behavior after September 2011 and on total time spent on the website, including the total number of webpage views and website visits by each applicant, though we do not have the latter information available by week. We have demographic information on age, education, race, and gender, as well as geographic location at the zip code level, for nearly all applicants. For vacancies, we have geographic location at the zip code level and the website's classification of each vacancy's industry. The industry classification, though, is more an amalgamation of industry and occupation (e.g., Management is an "industry" in our data). The industry classification is job-specific rather than firm-specific, so a sales position for a construction firm will be listed as part of Sales & Marketing rather than Construction.

We do not have information on the outcome of an application by a job seeker (i.e., contacted for an interview, hired, etc.), nor do we have any information on the job search efforts of a job seeker

outside of the website. We also do not know the employment status of a job seeker. Finally, we do not have information on when a vacancy was first posted to the website or when it was taken down, though we can identify the first and last day that any applicant on the website applied to a position within our sample period. In our analysis, we attempt to identify potential hires as instances when an individual's search spell ends with an application to a vacancy that disappears during the same week.

The raw data in our sample period contain over 46.15 million applicant-vacancy observations for 8.00 million job seekers and over 1.81 million vacancies. We remove matches with missing geography or industry data (a negligible amount of observations) as well as applications to postings that are to "work at home" jobs or similar postings that advertise positions that are self-employment opportunities rather than a job opportunity with a particular employer. This procedure reduces the number of individuals and vacancies in the data by a negligible amount, but it reduces the number of application observations by 14 percent. In some cases, we observe repeat applications by the same individual to the same job posting. We delete repeat posts that are within one week of each other on the presumption that these are the result of some form of an applicant error (i.e., multiple website clicks or forgetting that they just applied to the same job). We retain repeat applications further than one week apart on the presumption that applicants may find it worth it to apply again if they observe that the job opening is still posted (and, presumably, unfilled). We also exclude applicants with missing education and gender information, and missing website visit information (i.e., time spent on the website and number of webpage views), and we restrict our attention to applicants aged 16 to 75 years. These restrictions remove a negligible amount of individuals. We focus our analysis on the behavior of job seekers who register on the website after the beginning of our sample, September 1, 2010. This reduces our sample by about 31 percent but ensures that we only study behavior from the start of a job seeker's search spell. As we describe below, though, we use information on the behavior of all job seekers in our sample to generate estimates of vacancy characteristics, including total applications received.

2.B. Identifying Labor Markets and Active Vacancies

We define a job seeker's labor market by her metropolitan area. We restrict our sample to applicants within Core-Based Statistical Areas (CBSAs) that are not a Micropolitan Statistical Area, or Consolidated Statistical Areas (CSAs). The latter are combinations of two or more CBSAs that are considered part of a larger metropolitan area. The exclusion of Micropolitan Areas only removes the smallest metropolitan areas, defined as having a population between 10,000 and 50,000, and is done to ensure large enough sample sizes for our defined labor markets. This gives our analysis sample 318 metropolitan areas within the U.S. While we limit the scope of our analysis to job seekers within these areas, we keep track of all applications from these job seekers, regardless of the vacancy's location. Similarly, when calculating statistics for vacancies applied to by these job seekers, we use all applications to the vacancy, regardless of the origin of the application.

We use the labor market definition to identify the number of active vacancies present in a job seeker's labor market. We define a vacancy as active from the first date we observe a job seeker apply to the last date we observe someone apply. We refer to the time between the first application and the current period under analysis as the vacancy's duration. We refer to the time between the first and last observations as the vacancy's total tenure on the website. Obviously, our dating method creates a censoring issue for vacancies that were potentially active before or after our sample period. To account for this, we condition out calendar week means from all of the vacancy statistics that we generate.

2.C. Identifying Search Spells

Identifying the length and current duration of a job seeker's search spell is crucial to our analysis. The micro data provide information about when a job seeker registers with the search engine and when they are first and last observed applying to a vacancy during our sample period. The data do not provide the reason for why a job seeker stopped searching. A job seeker may have found

employment (either through the website or through other job search methods), stopped searching on the website, or stopped searching for a job altogether and left the labor force. It is important to note, too, that we have no information on the current employment status of the job seeker. Job seekers may search less intensely or intermittently because they are currently employed. They may also continue to search after finding employment (and likely at a different level of intensity) in cases where the new job reflects "underemployment" relative to the type of work they were seeking. We also do not have information on what other job search methods a job seeker is employing in addition to the search on the website. Consequently, our analysis only speaks to the search efforts a job seeker employs during her time on the website. We do make an effort to identify search spells that potentially end with a match to a vacancy on the website. We define these matches as spells that end with at least one application to a vacancy that ends the same week. We show below that job seekers are much more likely to apply to an ending vacancy at the end of a search spell regardless of their total spell length, so we believe that this proxy has suitable power for identifying potential matches.

A complicating issue in identifying job seekers' search spells is that many job seekers exhibit long spells of inactivity. Many job seekers in our sample spend only a short time (one week or less) on the website, either because they found work quickly or because they decided searching on the website did not suit their needs. There is a sizable fraction of job seekers, however, who reappear in our sample after a long spell of inactivity. We have no way to identify what these job seekers, whom we consider "marginally attached" to the website, were doing in during their period of inactivity. They may have left after finding a job but subsequently separated from the job after a short period. They may have become discouraged with the website and decided to employ other methods for a while, or they may have become discouraged with searching for work altogether. Finally, they may be either already employed or out of the labor force (with only a marginal desire for employment), and therefore searching intermittently and less intensely relative to an unemployed job seeker.

Taking these considerations into account, we define an individual's first search spell as the period between their first application date and any application date that is followed by more than five weeks of inactivity.[3] If we observe an individual submitting an application after the cutoff, we consider that the start of a new spell. We apply this rule throughout our sample. Note that the five-week rule is roughly consistent with the notion of unemployment in the Current Population Survey (CPS), which only counts someone as unemployed if they have done something to actively search for work within the past four weeks. In theory, the rule allows an individual to have up to seven distinct spells during our 53-week sample period. In practice, the majority of job seekers in our sample (66.4 percent) are on the website for one week. Many (40.3 percent of all job seekers in our sample) only apply to a single job. We identify only one search spell for most job seekers (82.9 percent) in our sample. For our main analysis, we restrict our attention to their first search spell. We do this because job seekers, given the range of reasons that generate multiple search spells, may not behave in the same way at the start of their search in subsequent spells. In our robustness exercises below, however, we show that the relationship between the duration of the spell and the number of applications sent during subsequent search spells is similar to behavior during the first spell, and we obtain similar regression results for these spells as we do for the first spell.

We also drop observations during the first week of our sample from our analysis because we cannot differentiate between new and incumbent vacancies during that week (a necessary criterion for testing stock-flow search behavior). We do, however, use the first week of data when calculating the aforementioned statistics. Finally, we aggregate most of our statistics for analysis to the weekly

[3] We also have information on the date that a job seeker registered with the website. Most job seekers (90.6 percent) send an application the same day they register for the website, and all but 5.1 percent send one within the first week.

interval.[4] This leaves us with an analysis sample consisting of 17.26 million applications to 1.41 million vacancies by 4.77 million job seekers. The sample also aggregates to 10.11 million job seeker-week observations, of which 7.65 million have at least one application sent that week.

2.D. Basic Evidence

Table 1 reports the demographic and socio-economic composition of our sample's job seekers during their first identified search spell on the website as compared with the composition of the unemployed and total labor force from the Current Population Survey (CPS). Our sample has a disproportional number of younger, minority, and less-educated job seekers relative to the labor force in the CPS. The demographic composition of our sample is closer to the demographic composition of the pool of unemployed, though it still over-represents the young and those with at least a college degree. A key difference between our sample and the pool of unemployed in the CPS is that our sample has a majority of female job seekers (56.9 percent) while in August 2011 the unemployed in the CPS are mostly male. In the appendix, we also compare the distribution of job seeker's search duration in our sample with the distribution of the unemployed in the CPS.

Table 2 shows that job seekers send an average of 1.7 applications per week. The average job search spell lasts just over two weeks. We find that job seekers apply to an average of 0.067 vacancies per webpage view, or about one application for every 15 webpages viewed. We view webpage views as a proxy for the number of vacancies job seekers peruse while searching on the website. Over the course of a completed spell, job seekers spend an average of 47 minutes per week searching on the website. Women send out slightly more applications per week than men, on average, but also send fewer applications per webpage view despite spending more time per week on the website. Older job seekers

[4] Job seekers tend to apply with at roughly seven-day intervals in the data, with large spells of inactivity and high volatility as to which day of the week individuals apply. For these reasons, we aggregate our data to the weekly frequency.

exhibit a similar pattern of website behavior, though with fewer applications relative to younger job seekers. One could interpret this as evidence that women and older workers are choosier in their job search. Such an interpretation would be consistent with women seeking more flexible employment, perhaps because of a disproportionate burden of childcare duties, and older workers who are more targeted in the types of jobs they seek, perhaps because of a greater accumulation of specific human capital. There is no discernible pattern across education groups or by race, though nonwhites appear to send somewhat more applications and spend somewhat more time per week on the website.

Table 3 lists the summary statistics of the characteristics of applications that were sent by job seekers in our sample. Just under 95 percent of a job seeker's applications are to vacancies within their own metropolitan area. Just under 76 percent are to a job seeker's modal industry, defined as the broad industry they applied to most often during their tenure on the website (including subsequent search spells after the first).[5] Just under 14 percent of a job seeker's applications are to a newly-posted vacancy (i.e., one that was posted the same week that she applied to it). Older, white, and more educated workers are more likely to apply to jobs outside of their metropolitan area, while older workers and women are more likely to apply to their modal industry. Men and older workers are more likely to apply to newer vacancies.

Figure 1 reports how the composition of applications in our sample changes with the duration of search. The upper left panel depicts the fraction of each week's applications that are to a job seeker's modal industry. The fraction falls with the duration of search. There is a sharp drop after the first week of search from 76.1 to 60.1 percent, though this drop is mostly due to the fact that job seekers who send only one application, by definition, send it to their modal industry. Among job seekers who search for at least two weeks, the fraction falls from 61.4 to 60.1 percent. After six months, only 55.0 percent of a job

[5] Just over 10 percent of job seekers have no modal industry.

seeker's applications are to her modal industry. We also examine application behavior conditional on demographics, completed spell length, and the number of active local vacancies (see Section 5 for the full specification). The conditional results in Figure 1 (as well as more detailed results in the appendix) suggest that the declining fraction of applications to one's modal industry is primarily attributable to long-duration job seekers searching broadly. The upper right panel of Figure 1 shows the fraction of applications to one's own metropolitan area, unconditionally and conditional on the above characteristics. In both cases, most applications (about 95 percent) are to one's own metropolitan area, regardless of the duration of search.

The lower left panel of Figure 1 shows that the fraction of applications sent to newly-posted vacancies rises by about 3 percentage points over the first six months of search but never rises above 17 percent through the first six months of search. Conditional on demographics, completed spell length, and the number of active local vacancies, however, the fraction is roughly constant over the duration of search. In the appendix, we show that controlling for the local stock of new and existing vacancies at each point during the search spell accounts for nearly the entire rise of the fraction of applications to newly-posted vacancies. Note, however, that the fraction sent to newly-posted vacancies is low: in the raw data, only about one in six applications are to a new vacancy even after six months of search. Finally, the lower right panel of Figure 1 shows that the fraction of "repeat" applications, i.e., applications to vacancies that a job seeker previously applied to, steadily increases over the duration of search. By the sixth month of search, nearly one in six applications are to a position that the job seeker has already applied to. Job seekers may interpret a vacancy that remains on the website as unfilled and send an additional application as a signal that they are still interested. The composition of job seekers explains about one-third of this rise: long-duration job seekers are more likely to send a repeat application throughout the duration of their search.

Finally, Table 4 reports the composition and basic characteristics of our sample's vacancies. Retail jobs are the most plentiful in our sample, accounting for 44.2 percent of all vacancies. They are followed by Food and Restaurant jobs (16.1 percent) and Customer Service jobs (8.9 percent). These shares are disproportionately higher than the employment shares of comparable occupations in the Occupational Employment Statistics Survey and of comparable industries in the Current Employment Statistics. The average vacancy received 31 applications and was active on the website for 6.5 weeks. There is large variation in the number of applications per week across industries, from 1.2 per week in Management positions to 9.5 per week in Warehousing positions.

3. Application Behavior by Duration

We now turn to the analysis of application behavior over the duration of search. Figure 2 shows the mean number of applications per week by weeks spent searching.[6] We report both the unconditional mean and applications per week after controlling for metropolitan area, the calendar week when a job seeker began her search, whether the week represented the end of a job seeker's spell, and whether the week represented the end of search during the sixth week of search. The latter two account for the fact that, by construction, a spell ends with at least one application, which produces a spike in the number of applications sent during the last week of search, and that spells that last exactly six weeks may have a disproportionately higher spike since they just fall within the bounds of inactivity that we use to define spell length. These controls do little to alter the pattern in the raw data, so we refer to them as our baseline specification. Job seekers send about 2.3 applications in their first week of search, on average. The number of applications falls to about 1.2 per week by the fourth week of search but then slowly starts rising. After six months of search, job seekers still send about 1.5 applications per week, on average.

[6] Throughout our analysis, we report confidence intervals based on robust standard errors that are clustered by metropolitan area.

We also estimate the effect of observable characteristics on application behavior by regressing applications for each duration week on our baseline model with additional controls for demographics (sex, a quadratic in age, and fixed effects for four education categories and four race categories) and whether the job seeker had any subsequent search spells. Figure 2 shows that there is little change in the basic relationship between application behavior and search duration.

Our data allow us to control for fixed unobservable job seeker characteristics. One way to do this is to control for a job seeker's completed spell length. The length of a completed search spell, to the extent that we can identify one in our data, is likely correlated with a variety of unobservable job seeker characteristics. Job seekers may differ in their search effort, leading low-search effort individuals to take longer to find work. Job seekers may also differ in their employability. Job seekers who are less desirable to prospective employers will take longer to find work. Job seekers may also differ in their search efficiency. Individuals who repeatedly apply to jobs that they are poorly suited for will take longer to find work. Finally, search spells may end in an outcome other than new employment for myriad reasons. For example, some job seekers may already be employed, and may exert nominal and infrequent search effort in the hopes of finding better employment. The same pattern may hold for individuals who are out of the labor force but entertain the idea of finding work. Other job seekers may simply give up on searching altogether. For example, Clark and Summers (1979) find that nearly half of all completed unemployment spells end with an exit from the labor force. Online job search has the added possibility that individuals continue searching for work, but not on the website. It is likely that a sizable fraction of the individuals whom we observe sending only one application are individuals who experimented with the website and decided it was not for them. It is unclear, though, how many individuals who stopped search in subsequent weeks came to the same conclusion.

Figure 3 depicts our baseline model from Figure 2 and compares it to a model where we add to the baseline model all of the job seeker characteristics described earlier and a set of fixed effects to control for spell length. Once we control for completed spell length, the declining relationship between applications and the duration of search steepens. It still exhibits a steep decline after the first week, from 2.9 to 0.8 applications per week, on average, but then continues to decline thereafter, to an average of just 0.1 application per week after six months of search.

The main advantage of our longitudinal data is that we can control for all fixed job seeker characteristics, both observable and unobservable, through the use of job seeker fixed effects.[7] Figure 3 shows a nearly identical decline in the number of applications per week, as compared to controlling for only completed spell length. This implies that completed spell length is a sufficient statistic for all fixed, unobservable job seeker characteristics relevant for the job seeker's application behavior.

We estimate the relationship between applications and duration separately for job seekers of differing completed spell lengths. Doing so sheds light on the behavior driving the declining application pattern in Figure 3. The results of this exercise are in Figure 4. We find that the number of applications sent per week by a job seeker declines as search continues and this pattern holds independently of the duration of the completed spell length. Regardless of a job seeker's completed spell length, applications per week decline steeply after the first week, then continue to decline throughout the subsequent weeks.

Another pattern emerges from Figure 4: individuals with longer spell lengths send more applications per week, on average, *throughout the duration of their search spell*.[8] The average number of applications per week is consistently monotonic in completed spell length. For example, in the third

[7] By construction, job seekers active for only one week contribute nothing to the identification in this case.
[8] Note that the figure excludes the last week of search because, by construction, every job seeker sends at least one application on the last week of their identified search spell.

week of search, those with completed spells of four weeks send an average of 0.8 applications, those with completed spells of 13 weeks send an average of 1.3 applications, while those with completed spells of 10 months or more send an average of 2.3 applications. The pattern holds throughout later weeks of search as well. The differences in application behavior by completed spell length are essentially a level effect. Thus, job seekers who search on the website longer send out more applications per week, on average, throughout their entire search spell.[9]

One may be concerned that the result is an artifact of how we define search spells. Our identification method requires that individuals apply for work at least once every five weeks, and imposes that the spell starts and ends with the first and last application observed, respectively. Two results, however, cut against such an interpretation. First, the mean number of applications per week for the longest-duration job seekers is well above two for most of their search. Individuals who were intermittently searching just enough to generate a long spell duration would likely have mean applications per week below one because a higher frequency of zero-application weeks would bring down the average. Second, the sheer magnitude of the difference in application behavior between long-duration and short-duration job seekers should lead to a higher probability of finding employment for the long-duration job seekers. Within the first month of search, individuals who search for 10 months or more send *triple* the number of applications as those who search for four weeks or less.

Nevertheless, we test the robustness of our result by replicating Figure 4 among the subset of job seeker-week observations where job seekers send at least one application. This eliminates the possibility of variations in the probability of zero-application weeks driving the results shown in Figure 4. The results of this exercise are in the left panel of Figure 5. The results show that the patterns in Figure 4 hold even when we eliminate observations when job seekers sent no applications during the week. We

[9] This pattern can be anticipated from Figure 3.

also estimate a version of Figure 4 that conditions out the variables used in our baseline specification as well the observable job seeker characteristics defined earlier and the weekly totals of new and existing vacancies defined in Section 5 below. The results are in the right panel of Figure 5. Again, the observed patterns from Figure 4 hold.

4. Application Behavior of Potential Matches

One may be concerned that the positive correlation between completed spell length and applications per week observed in Figures 4 and 5 (i.e., the finding that long-duration job seekers on average send more applications per period than the shorter-duration job seekers) arises only because individuals who did not find the website useful quit their search on it in favor of other methods. This can potentially generate a positive correlation between the total duration of search on the website and the number of applications sent that is not reflective of the job seekers' overall search effort.

To address this issue, we replicate the analysis of the previous section on the subset of individuals who potentially ended their search spell with a match from the website. This sample should minimize the effects of individuals who are only marginally interested in finding a job on the website.

4.A. Constructing the Potential Match Sample

To identify the subsample of the individuals who potentially found a job on the website, we restrict the subset to individuals who applied to at least one vacancy that expired during the same week as their search spell. The hypothesis is that, since the job seeker was one of the last people to apply to the vacancy, and since the vacancy was one of the last positions the job seeker applied to before ending their search spell, these individuals are the most likely to have exited her search spell by finding a job on

a website. We find that 22 percent of all job seekers in our sample satisfy this criterion, and 25 percent of job seekers who searched for at least two weeks satisfy the criterion.[10]

Figure 6 provides compelling evidence that this selection criterion provides a good identification of hiring through the website. The figure plots the average fraction of a job seeker's applications that are to an expiring vacancy each week, with job seekers grouped by their total spell length. We scale the horizontal axis so that it plots the remaining weeks of a job seeker's search spell. Week zero represents the last week of the spell. On average, 17 percent of all applications are to a vacancy that ends the same week that a job seeker applies to it. The fraction, however, increases sharply as job seekers reach the end of their search spell, with a substantial spike in the last week of search. The spike is considerably larger for job seekers with long search spells. Thus, we conclude that our selection criterion identifies potential matches between a job seeker and a vacancy. The subset of job seekers who apply to an expiring vacancy in week zero of Figure 6 constitutes our "potential match" sample.

4.B. Estimates Based on Potential Matches

Figure 7 plots the applications per week for the potential match sample. As one can see, the main patterns from Figure 4 are also present in Figure 7. In fact, the patterns are nearly identical.

We also replicate our regression-based estimates from Figures 3 and 4 using the potential match sample. Figure 9 shows the estimates from our baseline model and the baseline model extended to include controls for demographics and completed spell length. We report the estimates from using the subsample of job seekers with a potential match as well as the full sample of job seekers for comparison. The figure shows that the subsample yields qualitatively similar results. Quantitatively, the potential match sample shows a somewhat smaller effect of including completed spell length as a

[10] In comparison, Stevenson (2008) finds that 22 percent of all job seekers in 2002 found employment through the internet.

control, but there is still a monotonically declining relationship between applications per week and search duration.

4.C. Counterfactual Simulation of Potential Matches

One may worry that the potential match sample may still suffer from a spurious correlation between applications and search duration. Specifically, if our potential match sample contains a large number of individuals who are marginally interested in finding a job on the website (for example, because they have found a job on the website through a pure luck), one might worry that we obtain similar results in our full sample and potential match sample purely through this "luck" effect of stochastic job finding. Using a counterfactual simulation, we quantitatively evaluate such a possibility and show that pure luck in job finding cannot drive the results obtained from the potential match sample.

We do our simulation using a model of job seekers who only differ in their preference for search on the website. There are N total job seekers registered on the website. A fraction θ of these job seekers is what we refer to as "marginally attached" to the website. That is, they search both on the website and through other methods (including, potentially, other job search websites). We assume that $\theta = 0.8$, which is a fraction of "marginally attached" job seekers that is roughly consistent with the large amount of attrition we see within the first week of search.[11] All job seekers send n applications per week. To keep the exercise simple, we assume that the number of applications per week remains constant over the duration of search.[12] Job seekers who search exclusively on the website sent all n applications on the website. Those who are marginally attached send a fraction α of their applications through the website and the remaining $(1-\alpha)n$ applications to job openings found outside of the

[11] The exit hazard after one week of search is 74.3 percent.
[12] Constant search over time is assumed for simplicity given that this exercise focuses on behavior across individuals with differing completed spell lengths, and not differences within search spells.

website. In addition, marginally attached job seekers may quit the website entirely with probability $\rho(t)$, which we assume declines with search duration, t, given the sharp decline in job seekers observed in the data.[13] All job seekers have the same probability f of having an application lead to a hire each period, regardless of whether the application was made on the website or elsewhere.

Given the model setup, job seekers can exit search on the website in one of three ways: 1) they can find a job on the website, 2) they can find a job through other means, or 3) they can quit searching on the website entirely. Those who are marginally attached to the website can exit through any of the three methods, but those who are committed to the website can only exit through the first method. We do not allow job seekers to quit search entirely, however. They can only change their method of search over time.

The model has three parameters, $\{n, \alpha, f\}$ and one function, $\rho(t)$, that we calibrate to the data. We assume that n equals the mean number of applications sent in their first week of search by those applicants who completed spell lengths of at least 10 months. This is the highest amount of applications sent per week observed, on average, in the data, and is used since n represents the total number of applications sent using all methods in the model. We calibrate α using n and the model's expression for the expected total number of applications sent in the first week, $\theta\alpha n + (1-\theta)n$. We calibrate the job finding rate f to match the exit hazard of job seekers with completed spell lengths of six months or more. Given our assumption on the marginally attached, this exit hazard equals $1-(1-f)^n$. For the website quit probability, we assume that $\rho(t) = \rho_0/t^{\rho_0+1}$, which allows it to decline with duration as the p.d.f. of a Pareto distribution. We calibrate ρ_0 by equating the probability of exit after the first week to $1-(1-f)^n + \theta\rho(1)(1-f)^n$. We then run the model on 240,000 job seekers (roughly equivalent to 5 percent of our data sample), and use the results to generate the simulated versions of Figures 4 and

[13] We also perform the simulation under the assumption of a constant quit rate, and report these results in the appendix. The results are qualitatively similar to those we obtain with a declining quit rate.

7. In the simulated data, the potential match sample is the subset of job seekers who find a job through an application on the website. This sample will include those who were committed to search on the website and those who were marginally attached but managed to find a job through the website.

The results of the exercise are in Figure 8. The left panel shows the simulated applications per week for the full sample of simulated job seekers (analogous to Figure 4) and the right panel shows the simulated applications for those who found a job on the website (analogous to Figure 7). The simulation shows clear differences in applications per week by spell length between the full sample and the simulated sample. These differences are concentrated among the short-duration job seekers, who send much fewer applications per week than long-duration job seekers in the full sample, but nearly the same amount of applications per week in the potential match sample.

Intuitively, the marginally attached do not make up enough of the potential match sample to create much in the way of differences in application behavior (on the website) by completed spell length. Given our assumptions on the initial fraction of the marginally attached and an exponentially declining website quit rate, the marginally attached exit the website without finding a job and do so fairly quickly. This has two implications. First, relatively few of them find work on the website, leading to a small representation in the potential match sample. Second, many of them exit the website within the first few weeks (either through attrition or job finding elsewhere). Thus, they are concentrated within the short-duration job seekers. As a result, there is only a small difference in application behavior between the long-duration and short-duration job seekers within the potential match sample when the only thing that differentiates job seekers is their preference for search on the website. We could relax the assumption that the quit rate declines exponentially with duration, however, we would then not be able to match the empirical fraction of job seekers who exit the website within the first few weeks without assuming a counterfactually high job-finding rate for all job seekers. Furthermore, since the

marginally attached are also finding work outside of the website, it will still be the case that the marginally attached will make up a smaller share of the potential match sample, and the potential match sample will still have a less dispersed relationship between applications and duration than the full sample as a result.

Thus, neither "luck" in job finding of those who are only marginally interested in finding a job on the website, nor heterogeneity in preferences for the search on the website, can account for the finding that long-duration job seekers on average send more applications per week than the shorter-duration job seekers. Consequently, job seeker heterogeneity along other dimensions must drive the result that those with long spell lengths send more applications per week throughout their search.

5. The Role of Local Labor Demand

Before turning to additional exercises on robustness, we examine the role labor demand plays in accounting for the relationship between job seeker application behavior and search duration. Applications may decline over time simply because the pool of relevant vacancies shrinks over the duration of search. In stock-flow models of labor market search (Coles and Smith, 1998; Ebrahimy and Shimer, 2010), job seekers search over the full pool of available vacancies when they begin their search (i.e., the "stock"). If job seekers remain unemployed, they subsequently search only over newly-posted vacancies as they arrive (i.e., the "flow"). Consequently, these models predict that applications per week should drop precipitously after the first week of search and remain roughly constant thereafter (assuming a roughly constant arrival rate of new vacancies). In the online application data we also observe a discontinuous drop in the number of applications after the first week of search and a steady decline afterwards. Since we have detailed data on the timing of applications to all vacancies on the website, we can construct measures of the stocks of total and newly-posted vacancies within each metropolitan area during each calendar week of the sample. These measures allow us to estimate the

effect that any stock-flow process of job search may have on the relation between application behavior and search duration.

We estimate our baseline model from Figure 1, but add to it a flexible formulation of new and incumbent vacancies as in the following model:

$$(1) \qquad A_{ij\tau}(t) = d(t) + \alpha_j + \gamma_\tau + \eta_T(t) + \beta_0 \ln v_j^n(t) + \beta_1 \ln v_j^e(t) + \beta_2 d(1) \cdot \ln v_j^e(t) + \varepsilon_{ij\tau}(t).$$

For individual i in metropolitan area j during duration week t, the model regresses total applications during the week, $A_{ij\tau}(t)$, on a set of dummy variables for the week of search, $d(t)$. It includes fixed effects for metropolitan area j and spell start week τ, as well as the control for whether the week represents the last week of the spell, from the baseline model. It also includes the (log) number of vacancies posted during the week, $v_j^n(t)$, the (log) number of incumbent vacancies that existed prior to the week and remain active, $v_j^e(t)$, and the (log) number of incumbent vacancies interacted with an indicator for the first week of search. A stock-flow model would imply that $\hat{\beta}_1 = 0$, while a model where the full stock of vacancies is sampled each period would imply that $\hat{\beta}_2 = 0$.

Our estimates from this model are in Figure 10. We estimate a model where we use completed spell length as a control for unobservable job seeker characteristics, and a model where we use job seeker fixed effects. In either specification, controlling for both the number of active local vacancies does little to alter the baseline relationship between applications per week and search duration. The predicted number of applications per week is slightly higher when controlling for job seeker fixed effects rather than observable job seeker characteristics and spell length fixed effects, but the basic pattern holds—a job seeker's search effort, as measured by applications per week, falls with the duration of job search.

Given the evidence in Figure 1, i.e., that the majority of a job seeker's applications continue to be to existing rather than to new vacancies months after the search spell begins, it is not surprising that controlling for the stock of new and existing vacancies does little to affect the relationship between application behavior and search duration.. Even though the fraction of applications to new vacancies rises with duration, consistent with a stock-flow search process, it never rises to a level that is quantitatively large enough to affect the strong, declining relationship between applications and search duration.

6. Robustness

To further reinforce the view that our findings reflect differences in search behavior across job seekers and over time, and are not artifacts of examining job search on a single online search engine, we present additional results as checks of robustness. First, we examine search behavior during the second and subsequent search spells for job seekers whom we identify as returning to the website after a period of inactivity on the website. This exercise addresses a concern that the decline in the number of applications over the course of search reflects more efficient search by the job seekers who have learned how to successfully navigate the website over time. Second, we replicate our results for strong and weak labor markets. This exercise sheds more light on our finding that long-duration job seekers on average send more applications in every period of job search. If this is due to an increased effort that compensates for poor expectations on job prospects, then job seekers in weak labor markets, all else equal, should send more applications per week. Finally, we examine evidence on total search time spent on the website and on the number of applications per website view by completed spell length. We only have information available on the total website time and views for each job seeker, but we can still use this information to shed more light on what drives our main results. The number of applications per website view can be thought of as a proxy for "choosiness" of job seekers.

6.A. Evidence from Multiple Search Spells

We examine whether the second and subsequent search spells on the website, identified using the five-week cutoff, exhibit qualitatively similar application behavior as the one documented for the first search spell after registration on the website. In doing so, we identify job seekers with two or more spells and stack the job seeker-week observations of these spells with the first-spell observations of our main sample. We then replicate our previous regression analyses on the stacked panel, including dummy variables for the spell number and interactions between the spell number and the current duration of the spell. We identify a second spell for about 17.3 percent of job seekers, a third spell for 4.0 percent of job seekers, and a fourth or higher spell for about 0.9 percent of job seekers. In the regression analysis, we use a single dummy variable for the fourth and subsequent spells because of the relatively small sample size for this group of job seekers and the fact that later spells are increasingly right-censored given the one-year length of our sample period.

Figure 11 presents the results of this exercise.[14] We replicate our regression analysis using our baseline specification and the full specification that includes additional controls for jobseeker fixed effects and the number of incumbent and newly-posted vacancies active in the metropolitan area. The figure shows that the later search spells all exhibit a declining number of applications per week over their duration. In fact, their patterns are nearly identical to those one observes for the first spell. Thus, the evidence rejects the hypothesis that the decline in applications per week may reflect more efficient search due to learning about the website, and instead it supports the hypothesis that search effort declines with the duration of a search spell.

6.B. Application Behavior and Labor Market Tightness

[14] We report the estimates for the first three spells given the noisy nature of the estimates for the fourth and subsequent spells.

We also examine application behavior in strong versus weak labor markets. We do so by splitting the metropolitan areas in our sample into quartiles based on two measures of labor market strength. The first is the CBSA unemployment rate, taken from the BLS Local Area Unemployment Statistics (LAUS) data, and averaged over September 2010 through September 2011. The second is a measure of labor market tightness derived from the website data. It is the ratio of active vacancies to active job seekers, averaged across all weeks in the sample period. A higher ratio represents a stronger (i.e., tighter) labor market. We focus on comparisons of CBSAs in the highest and lowest quartile based on each measure. The LAUS unemployment measure is a broad measure of labor market strength, while the market tightness measure is specific to job seekers and vacancies on the website.

Our estimates are in Figure 12. We report the results using the full specification of our model that controls for completed spell length (i.e., a replication the specification in Figure 10). The left panel reports the results for high-unemployment and low-unemployment CBSAs while the right panel reports the results for high-market tightness and low-market tightness CBSAs. In both cases, job seekers in weak labor markets send more applications per week throughout the duration of their search spell. The difference in applications per week is somewhat small in absolute value, ranging between 0.08 and 0.18 applications per week (between 4 percent and 12 percent of the sample mean), depending on the duration of search, but it is statistically significant throughout most of the search spell. Our findings are consistent with the notion that job seekers exert more effort when they perceive relatively grim employment prospects. It supports an interpretation of the evidence in Figures 4, 5, and 7 where long-duration job seekers exert the most effort because they anticipate a low probability of finding work.

6.C. Search Time and Application Selectivity

Finally, we examine the evidence we have on website search behavior outside of sending an application. We have data on the total time each job seeker spent on the website and their total number

of webpage views, but we only have this information in the aggregate, not broken out by week. Therefore, we examine these data by completed spell length rather than duration of the spell. Doing so implies that we cannot distinguish between differences due to ex ante heterogeneity and differences due to changes in behavior over the duration of search. Thus, the evidence should be viewed as complementary to our main analysis rather than a robustness check on its validity.

We focus on two statistics. The first is the number of applications sent per webpage view, measured as total applications sent during the job seeker's total time on the website (across all search spells) divided by total webpage views over the same period. We consider the measure a proxy for the "choosiness" of the job seeker in their application behavior, under the assumption that webpage views are a proxy for the number of job listings the job seeker reads through before applying for a job. A higher ratio implies less choosiness. The second statistic is the average search time per week spent on the website, measured as total time spent browsing the website divided by total tenure on the website (across all search spells). This statistic serves as an alternative measure of search effort.

Figure 13 presents the results. The left panel of the figure shows that the number of applications per webpage view is increasing in completed spell length. Job seekers who search for up to a month send an application for every 15.8 webpages viewed, on average, while those who search for six months send an application for every 13.8 webpages viewed, on average. We interpret this as evidence that long-duration job seekers are less "choosy" in the vacancies that they apply to, though we cannot distinguish whether they are less choosy throughout their search, or whether choosiness declines with duration. Evidence on application behavior from Figure 1, however, suggests that they may be less choosy throughout their search, since completed spell length accounts for most of the rise in the fraction of applications sent outside one's modal industry and the fraction of repeat applications.

The right panel of Figure 13 shows that time spent searching per week declines sharply with completed spell length and remains relatively stable for most individuals with longer search spells. Individuals with spells less than a month in length spend over 50 minutes per week on the website, while those with completed spells of six weeks or longer spend just over 25 minutes per week on the website. Again, we cannot differentiate whether long-duration job seekers spend less time searching regardless of spell length or whether time spent searching declines with duration; however, in the appendix, we show that applications per week exhibit a similar pattern to search time per week when calculated by completed spell length. Thus, we consider the number of applications per week to be a suitable measure of search effort on the website.

7. Conclusions

This paper examines the relationship between job seekers' search behavior and the duration of search using data from an online job search website. We focus on the number of applications job seekers send to vacancies posted on the website as our measure of search effort.

We find that a job seeker's search effort declines with search duration. Our findings demonstrate that unobserved job seeker heterogeneity is an important part of behavior over the course of job search. In particular, if we do not control for fixed job seeker characteristics, the average number of applications sent each period from the start of the search initially declines but then starts increasing. This increase reflects a composition effect: job seekers with longer completed spell lengths generally send more applications per period throughout the duration of their search.

These findings are robust to analyses that account for various measurement issues, including infrequent job search and unknown reasons for the end of a search spell. The results also carry through within a subsample of job seekers whom we identify as generating a potential match with a vacancy on

the website. We also control for the number of active vacancies in a job seeker's metropolitan area, differentiating between pre-existing and newly posted vacancies. We find that it has little effect on the application-duration relationship. This is because most applications are sent to pre-existing vacancies, and there is little change in the stock of these vacancies over the duration of search. Thus, controlling for both the composition of job seekers and local labor demand produces a strongly declining relationship between applications per week and search duration.

Our finding that long-duration job seekers send the most applications per week throughout their entire duration of search points to the importance of job seeker heterogeneity. Ignoring this heterogeneity would make this finding appear at odds with theories in which high effort job seekers find jobs faster. However, we interpret the finding as evidence that unobserved heterogeneity in job seeker characteristics is quantitatively more important than differences in search effort for explaining why some job seekers have long search spells. This heterogeneity maybe due to long-duration job seekers being among the least employable or the least efficient in targeting their search toward vacancies for which they are well-suited.

References

Addison, John T., and Pedro Portugal, 1989. “Job Displacement, Relative Wage Changes, and Duration of Unemployment,” *Journal of Labor Economics*, Vol. 7(3): 281-302.

Barron, John M., John Bishop, and William C. Dunkelberg, 1985. “Employer Search: The Interviewing and Hiring of New Employees,” *Review of Economics and Statistics,* Vol. 67(1): 43-52.

Blanchard, Olivier Jean, and Peter Diamond. 1994. "Ranking, Unemployment Duration, and Wages." *Review of Economic Studies*, Vol.61(3): 417-434.

Brown, Jennifer, and David Matsa, 2012. “Boarding a Sinking Ship? An Investigation to Job Applications to Distressed Firms,” mimeo, Kellogg School of Business, Northwestern University.

Burdett, Kenneth and Cunningham, Elizabeth, 1998. “Toward a Theory of Vacancies,” *Journal of Labor Economics*, 16 (3): 445–478.

Clark, Kim B., and Lawrence H. Summers, 1979. “ Labor Market Dynamics and Unemployment: A Reconsideration.” *Brookings Papers on Economic Activity*, Vol. 1979 (1): 13-72.

Coles, Melvyn G., and Eric Smith, 1998. "Marketplaces and Matching," *International Economic Review*, Vol. 39(1): 239-54.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger, 2013. “The Establishment-Level Behavior of Vacancies and Hiring,” *Quarterly Journal of Economics,* forthcoming.

Ebrahimy, Ehsan, and Robert Shimer, 2010. “Stock-Flow Matching,” *Journal of Economic Theory*, Vol. 145(4): 1325–1353.

Holzer, Harry J., Lawrence F. Katz, and Alan B. Krueger, 1991. “Job Queues and Wages,” *Quarterly Journal of Economics*, Vol. 106(3): 739-768.

Hornstein, Andreas. 2012. “Accounting for Unemployment: The Long and Short of It”, The Fderal Reserve Bank of Richmond, mimeo.

Jones, Stephen R. G., 1988. “The Relationship between Unemployment Spells and Reservation Wages as a Test of Search Theory.” *Quarterly Journal of Economics*, Vol. 103 (4): 741-765.

Kuhn, Peter,and Mikal Skuterud, 2004. “Internet Job Search and Unemployment Durations,” *American Economic Review*, Vol. 94(1): 218-232.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo, 2012. “Duration Dependence and Labor Market Conditions: Theory and Evidence from a Field Experiment,”NBER Working Paper No. 18387.

Kroft, Kory, and Devin Pope. 2012. “Does Online Search Crowd-out Traditional Search and Improve Matching Efficiency? Evidence from Craigslist“. *Journal of Labor Economics*, Forthcoming.

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 (1): 1-81.

Kudlyak, Marianna, Damba Lkhagvasuren, and Roman Sysuyev, 2012. “Sorting by Skill over the Course of Job Search,” Federal Reserve Bank of Richmond Working Paper No. 12-03.

Lockwood, Ben. 1991. "Information Externalities in the Labour Market and the Duration of Unemployment." *Review of Economic Studies*, Vol. 58(4): 733-753.

Machin, Stephen, and Alan Manning. 1999. The Causes and Consequences of Long-term Unemployment in Europe. in: O. Ashenfelter and D. Card (ed.), Handbook of Labor Economics, Vol. 3, Ch. 47: 3085-3139.

Marinescu, Ioana, and Ronald Wolthoff, 2012. “Wages, Applications, and Skills,” mimeo, Department of Economics, University of Toronto.

Mortensen, Dale. T.. 1977. "Unemployment Insurance and the Labor Supply Decision," in: A. Katz, ed., "The economics of unemployment insurance: a symposium", *Industrial and Labor Relations Review*, 30.

Moscarini, Giuseppe. 2001. "Excess Worker Reallocation," *Review of Economic Studies*, 69(3): 593-612.

Ours, Jan van and Geert Ridder. 1991. “Cyclical Variation in Vacancy Durations and Vacancy Flows.” *European Economic Review,* 35, 1143-1155.

Pissarides, Christopher A.. 1992. "Loss of Skill During Unemployment and the Persistence of Employment Shocks." *Quarterly Journal of Economics*, 107(4): 1371-1391.

Pissarides, Christopher A., 2000. *Equilibrium Unemployment Theory*, 2nd Edition, MIT Press Books: The MIT Press.

Rogerson, Richard, Robert Shimer, and Randall Wright. 2005. "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, Vol. 43: 959-988.

Stevenson, Betsey, 2008. “The Internet and Job Search,” NBER Working Paper No. 13886.

van den Berg, Gerard, and Jan van Ours. 1996. "Unemployment Dynamics and Duration Dependence", *Journal of Labor Economics*, 14(1).

Table 1. Demographic Characteristics, Website Sample and the Current Population Survey

	Share of Website Job seekers	Share of Unemployed (CPS)	Share of Labor Force (CPS)
Gender			
Male	43.1	56.3	53.3
Female	56.9	43.7	46.7
Age			
16-24 Years Old	52.8	26.3	13.6
25-39 Years Old	26.9	31.6	32.2
40-54 Years Old	15.2	27.4	34.2
55+ Years Old	5.1	14.7	19.9
Education[1]			
High School or Less	50.7	51.0	37.1
Certification or Some College	15.0	19.5	17.1
Associates Degree	14.8	20.0	10.6
Bachelors Degree or More	19.5	9.4	35.1
Race			
White	50.3	54.4	67.2
Black	25.4	19.4	11.0
Hispanic	14.6	19.2	14.8
Other	9.7	6.9	6.9

Notes: Table reports the share of individuals in each demographic category from our sample of website job seekers as well as the unemployed and those in the labor force, as reported in the Current Population Survey (CPS). CPS statistics are monthly averages over September 2010 to September 2011.

Table 2. Search Behavior by Demographic Group

	Applications per Week	Total Weeks Spent Searching	Applications per Webpage View	Search Time per Week (mins)
Full Sample	1.71	2.13	0.067	47.0
Male	1.67	2.13	0.071	44.1
Female	1.74	2.12	0.065	49.2
16-24 Years Old	1.88	1.97	0.066	48.0
25-39 Years Old	1.66	2.06	0.073	43.4
40-54 Years Old	1.44	2.49	0.064	47.8
55+ Years Old	1.33	2.94	0.058	53.6
High School Degree	1.74	2.04	0.068	46.8
Certification	1.62	2.27	0.068	48.0
Associates Degree	1.70	2.39	0.065	48.7
Bachelors Degree	1.69	2.21	0.066	46.5
Graduate Degree	1.48	2.03	0.072	40.2
White	1.65	2.09	0.067	44.0
Black	1.78	2.20	0.068	50.6
Hispanic	1.76	2.06	0.066	49.4
Other	1.72	2.13	0.067	48.4

Notes: Summary statistics are mean values across all job seekers in our website sample. Applications per week and total weeks spent searching are for the first identified search spell only. Applications per webpage view and search time per week are for the job seeker's full tenure on the website.

Table 3. Characteristics of Vacancies Applied to by Demographic Group

	Pct. to Same CBSA	Pct. to Modal Industry	Pct. to Newly-Posted Vacancy
Full Sample	94.7	75.8	13.7
Male	94.5	74.7	14.5
Female	94.9	76.6	13.0
18-24 Years Old	95.6	75.9	10.8
25-39 Years Old	93.9	75.5	15.7
40-54 Years Old	93.5	75.7	18.2
55+ Years Old	93.6	76.6	19.3
High School Deg.	95.6	76.4	12.1
Certification	93.8	74.5	16.4
Associates Deg.	94.2	73.8	15.2
Bachelors Deg.	92.5	75.4	16.1
Graduate Deg.	89.7	79.2	22.4
White	93.6	76.3	14.0
Black	96.0	74.5	12.6
Hispanic	96.4	76.5	13.5
Other	95.0	76.2	15.2

Notes: The table reports the fraction of all applications made by job seekers in our website sample during their first identified search spell for each category. Modal industry is defined as the broad industry category that job seekers sent the most applications to (and is the industry of the sole application if the job seeker only applied to one job on the website). Newly-posted vacancies are those that first appeared on the website within a week of the job seeker's application to it.

Table 4. Characteristics of Vacancies by Industry

	Share of Vacancies	Mean Total Applicants	Mean Web Tenure (weeks)	Applicants per Week Active
Full Sample	100.0	31.1	6.5	4.8
Personal Care	2.0	75.5	8.3	9.1
Other Services	0.8	57.8	13.8	4.2
Hotel & Hospitality	2.2	43.9	5.4	8.2
Food & Restaurant	16.2	38.5	9.6	4.0
Education	0.7	35.8	23.2	1.5
Retail	44.6	34.5	5.8	5.9
Automotive	2.6	33.1	17.6	1.9
Transportation	1.3	32.2	4.3	7.6
Warehousing	2.8	29.2	3.1	9.5
Administrative & Office	1.4	27.1	2.9	9.4
Construction	2.0	22.0	6.7	3.3
Professional & Technical	4.7	17.8	2.5	7.1
Management	2.0	16.4	14.1	1.2
Sales & Marketing	3.3	15.7	4.2	3.8
Health & Wellness	4.7	13.3	5.1	2.6
Customer Service	8.8	12.4	2.6	4.8

Notes: Summary statistics are for all vacancies applied to by job seekers during their first search spell in our website sample. “Mean total applicants” refers to the average during their vacancy’s total tenure on the website. “Mean web tenure” is the average total time between the first and last application to a vacancy on the website. Both measures condition out calendar week effects.

Figure 1. Application Behavior over the Duration of Search

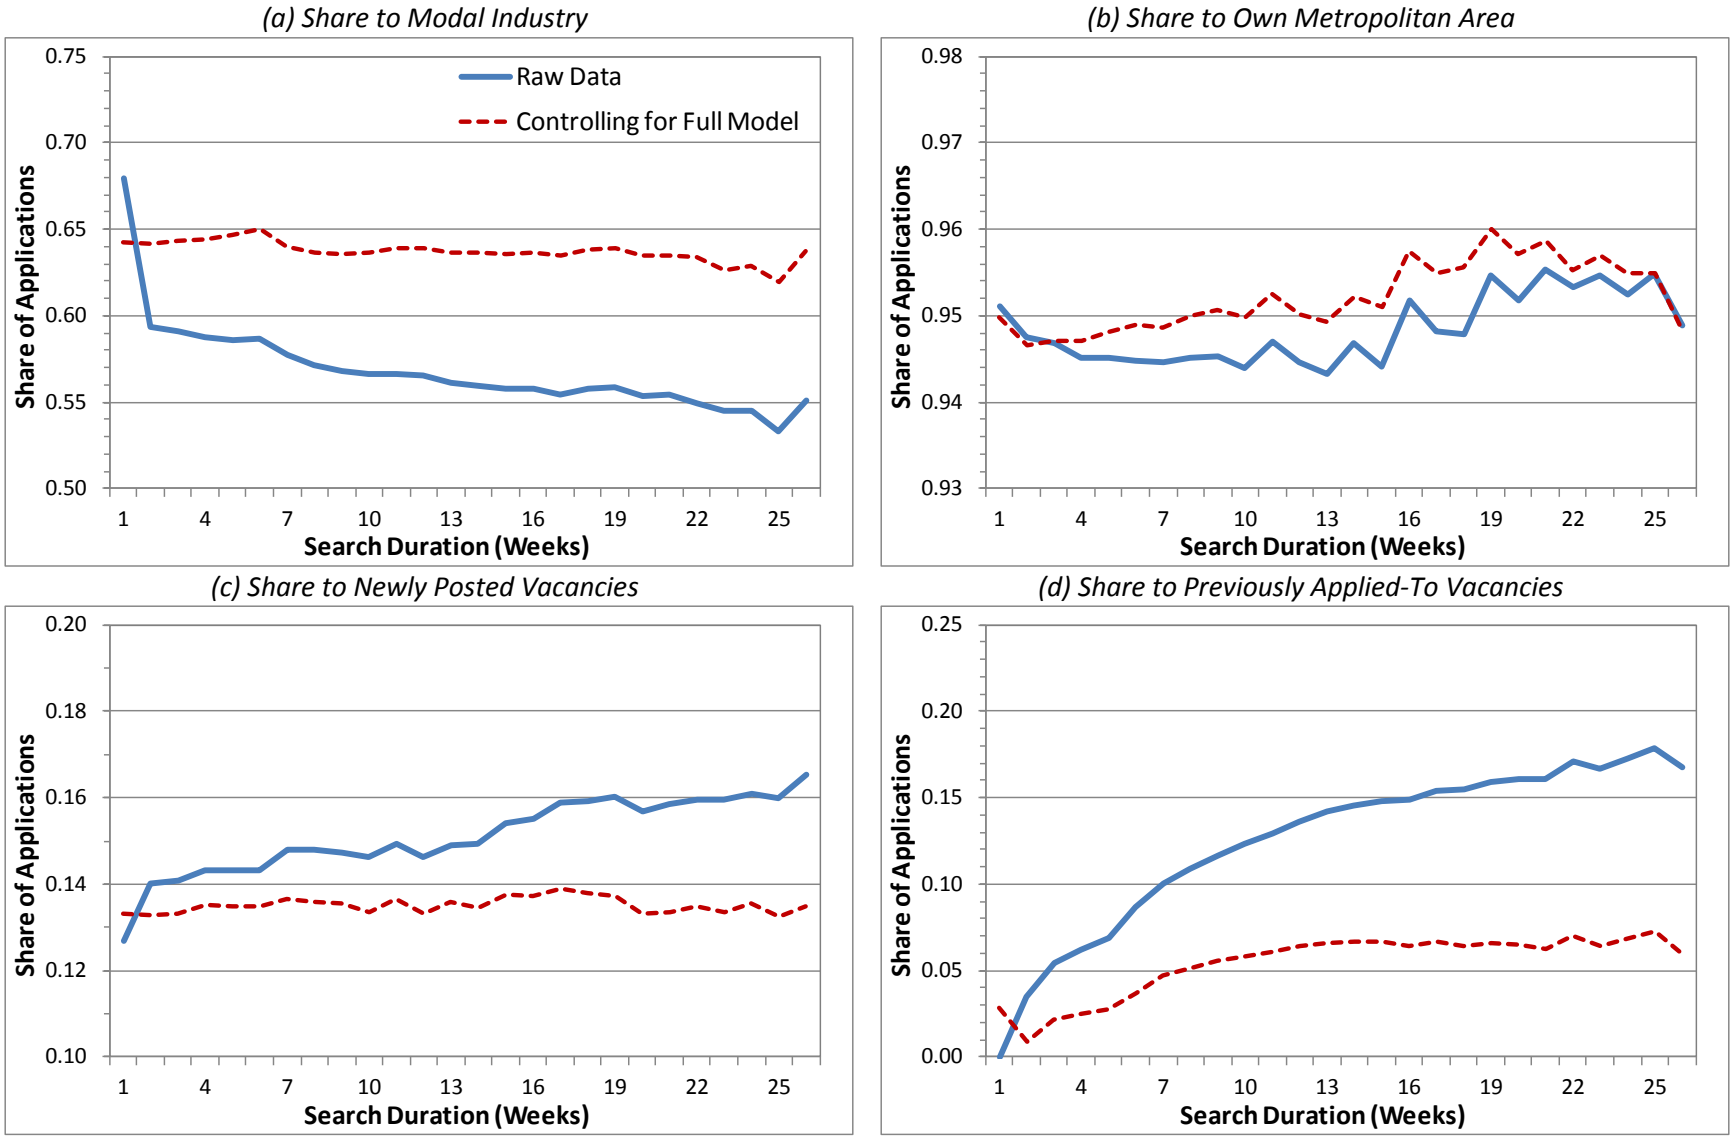

Notes: Panels depict the (unconditional) fraction of applications to a job seeker's modal industry, to the job seeker's metropolitan area, to newly-posted vacancies, and to previously applied-to vacancies by search duration, separately by the total spell length of job seekers.

Figure 2. Applications over the Duration of Search, Controlling for Demographics

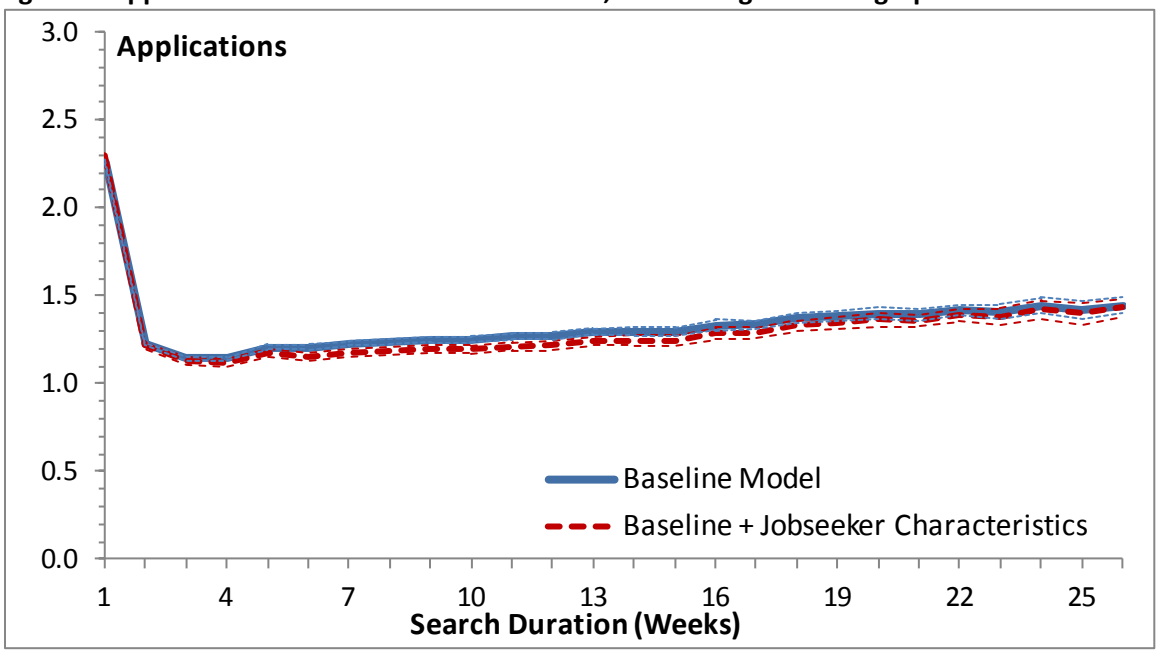

Notes: Figure shows estimated relation between applications per week and duration of search in our baseline model (which controls for the job seeker's metropolitan area and the start and end of the job seeker's spell) and a model that additionally controls for demographic characteristics (age, age squared, gender, education, and race). Dashed lines represent 95 percent confidence intervals (with standard errors clustered on job seeker's metropolitan area).

Figure 3. Applications over the Duration of Search, Controlling for Fixed Job seeker Characteristics

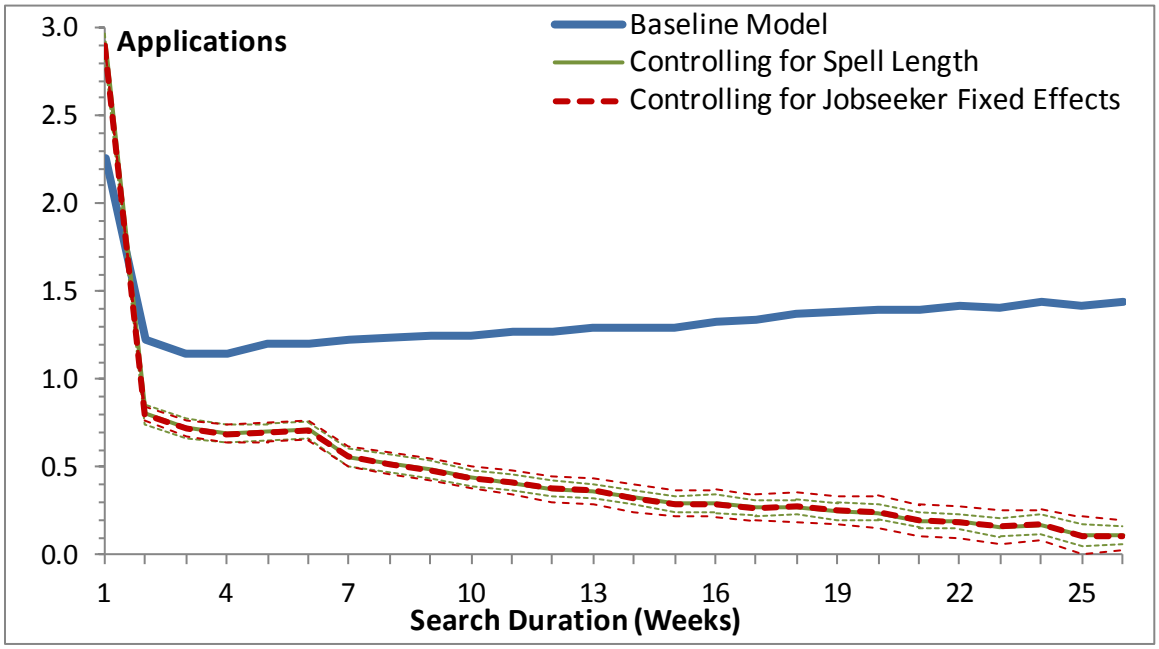

Notes: Figure shows the estimated relationship between applications per week and the duration of search for our baseline model, a model that additionally controls for total spell length, and a model that includes person fixed effects. Dashed lines represent 95 percent confidence intervals (with standard errors clustered on job seeker's metropolitan area).

Figure 4. Applications over the Duration of Search by Completed Spell Length

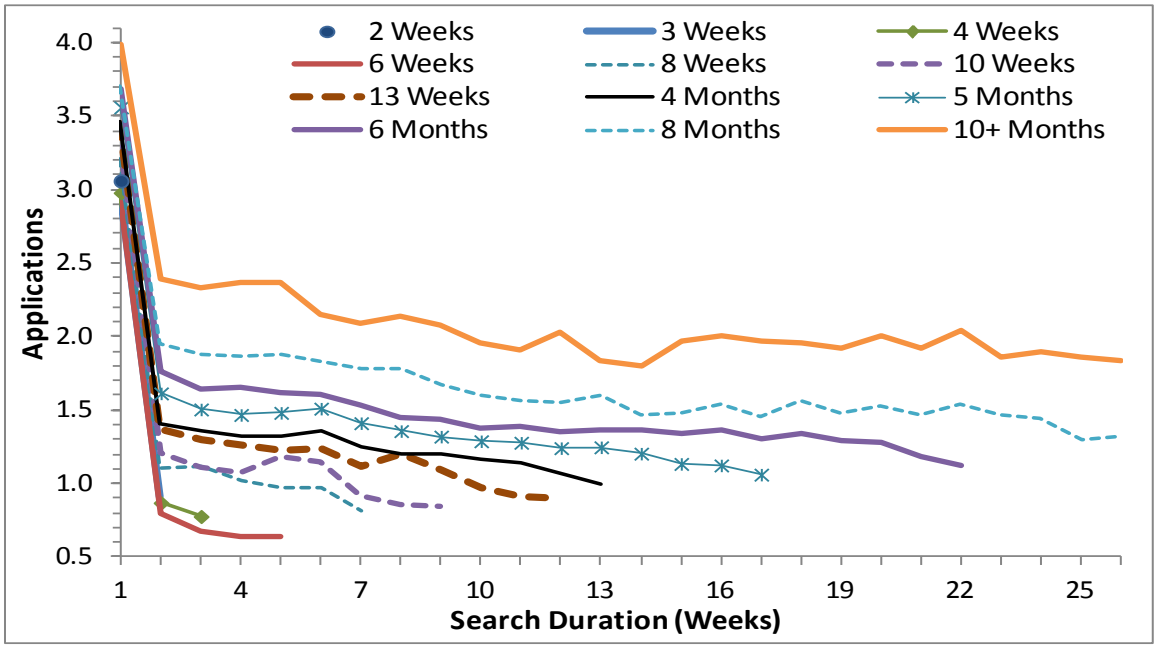

Notes: Figure shows the estimated (unconditional) relationship between applications per week and duration of search separately for job seekers based on the total length of their search spell. Only selected spell lengths are reported.

Figure 5. Applications over the Duration of Search by Completed Spell Length,

(a) Conditional on Sending at Least One Application

(b) Conditional on Observable Worker and Labor Market Characteristics

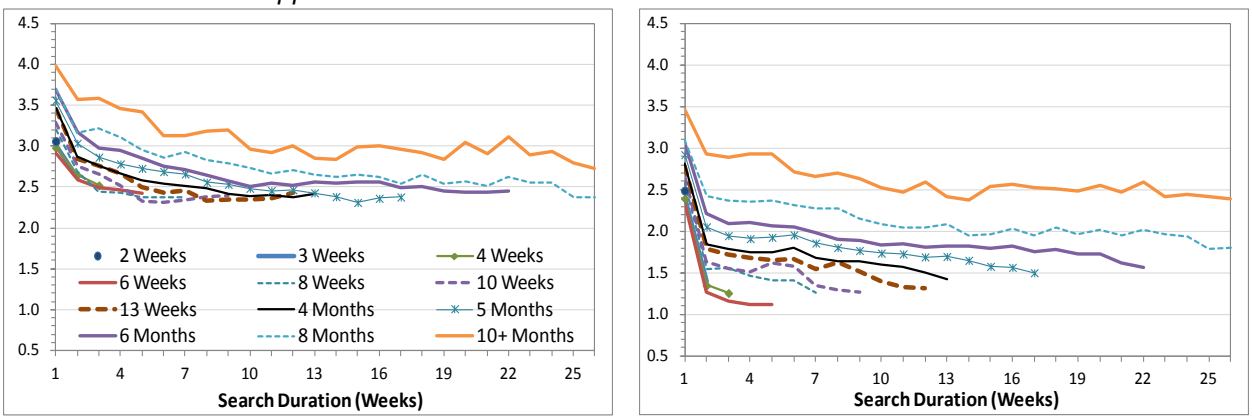

Notes: Figure shows the estimated (unconditional) relationship between applications per week and duration of search separately for job seekers based on the total length of their search spell. In the left panel, mean applications are only calculated for individuals who sent at least one application in a given week. In the right panel, mean applications are calculated for all individuals, but only after conditioning out demographic and local labor market characteristics. See text for details. Only selected spell lengths are reported.

Figure 6. Fraction of Applications to Expiring Vacancies by over the Duration of Search by Completed Spell Length

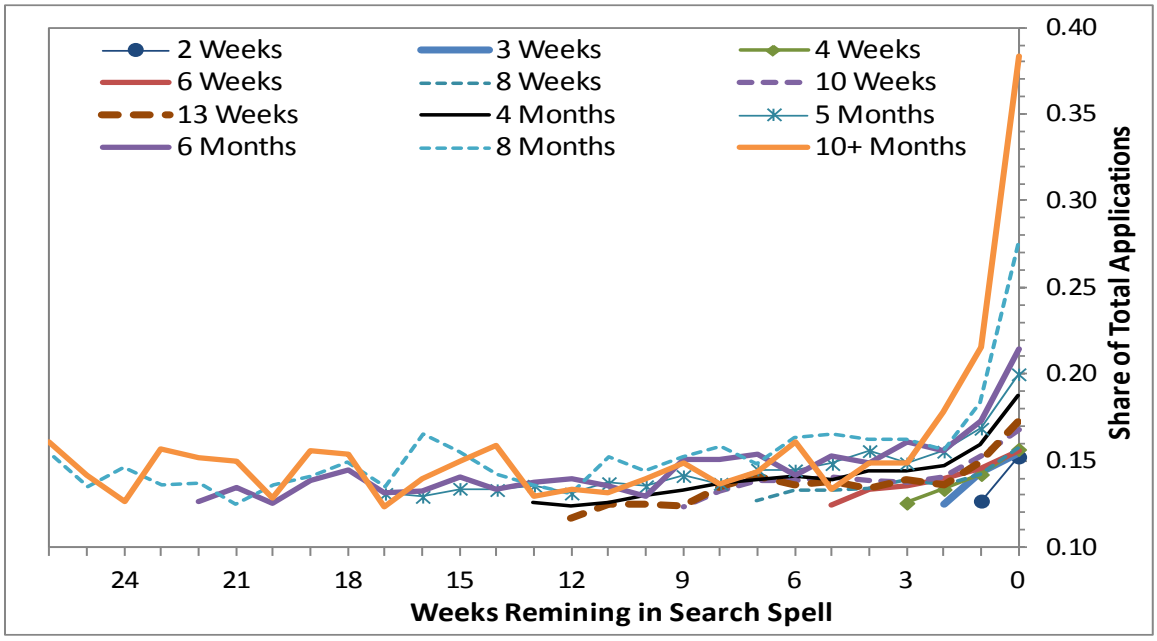

Notes: Figure shows the estimated (unconditional) fraction of applications each week that are to an expiring vacancy (defined as a vacancy that receives its last application during the same week) by the number of weeks remaining in the search spell, separately by the total spell length of the job seeker. Only selected spell lengths are reported.

Figure 7. Applications by Search Duration and Completed Spell Length, Conditional on Ending Search with an Application to an Expiring Vacancy

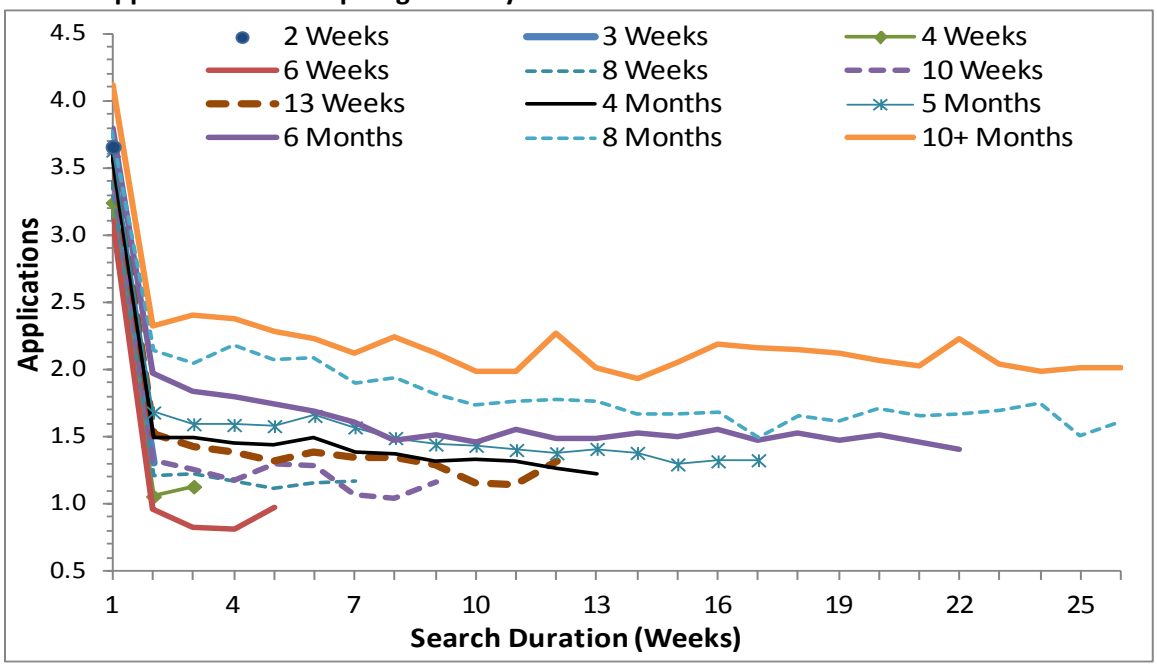

Notes: Figure shows the estimated (unconditional) relationship between applications per week and duration of search separately for job seekers based on the total length of their search spell. Mean applications are only

calculated for individuals who sent at least one application to a vacancy that expired (defined as receiving its last application) during their last week of search. Only selected spell lengths are reported.

Figure 8. Applications over the Duration of Search, Various Controls, Potential Match Sample

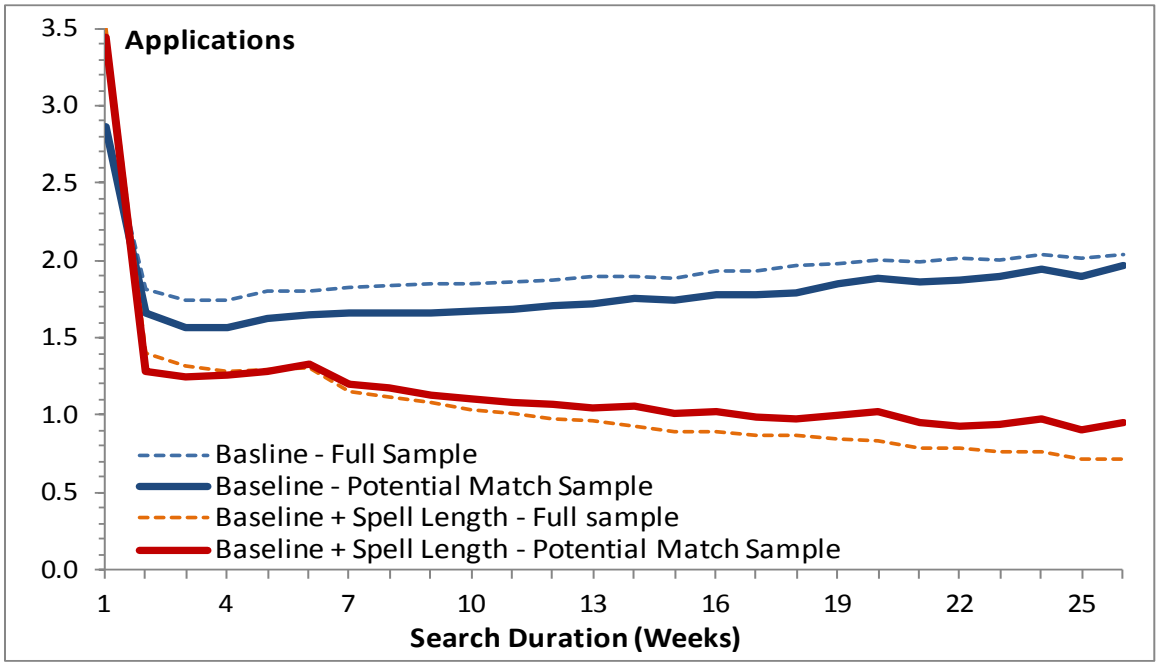

Notes: Figure shows the estimated relationship between applications per week and the duration of search for our baseline model and a model that additionally controls for total spell length using the full sample of job seekers and a subsample that are identified as potentially matching with an expiring vacancy.

Figure 9. Simulated Application Behavior by Completed Spell Length, Heterogeneous Tastes for Website Search

(a) Full Sample

(b) Simulated Match Sample

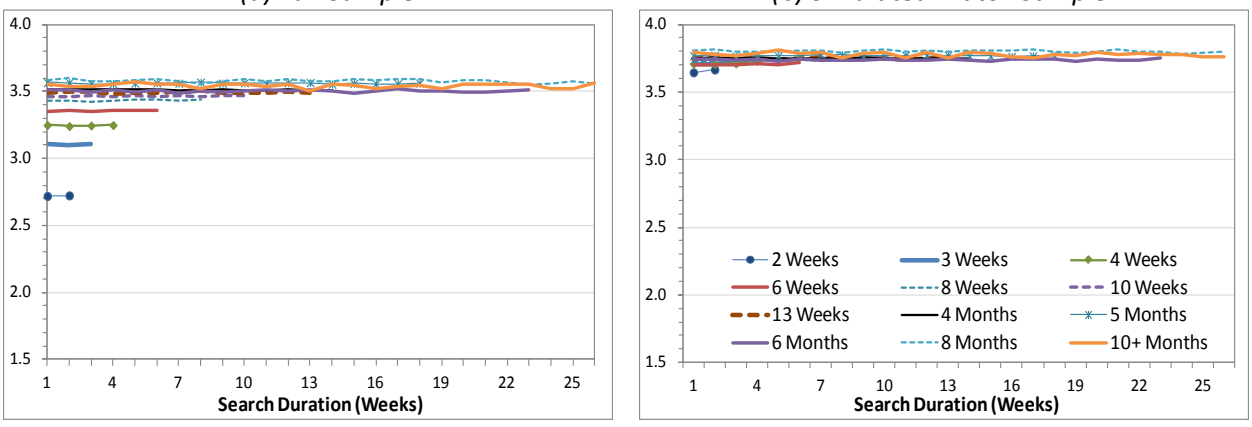

Notes: Figure shows the estimated (unconditional) relationship between applications per week and duration of search separately for job seekers based on the total length of their search spell using a simulated sample of job seekers calibrated to the empirical distribution of job seekers in our website sample. The left panel reports the estimates for all simulated job seekers, while the right panel reports the estimates for simulated job seekers who found employment through the website. Only selected spell lengths are reported.

Figure 10. Applications over the Duration of Search, Controlling for Active Vacancies and Fixed Job seeker Characteristics

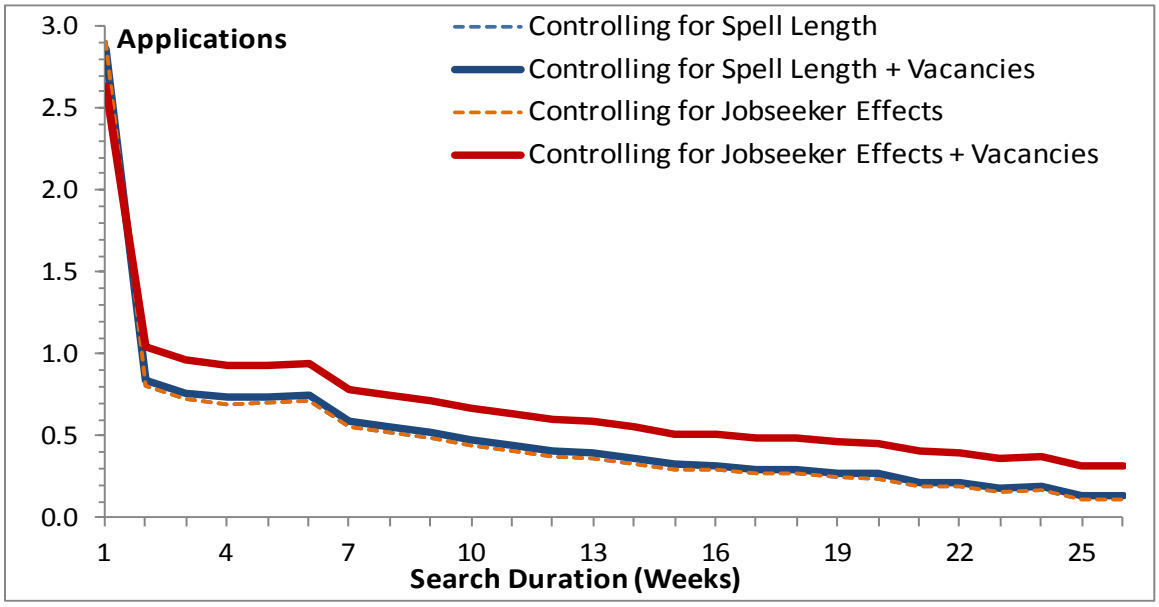

Notes: Figure shows the estimated relationship between applications per week and the duration of search for our baseline model, a model that additionally controls for the number of active vacancies, and a model that controls for active vacancies and fixed job seeker characteristics. The latter are controlled for by either fixed effects for spell length and demographic controls or job seeker fixed effects.

Figure 11. Applications over the Duration of Search, Estimated with Multiple Spells per Job seeker

(a) Baseline Model *(b) Controlling for Spell Length, Job seeker Effects, and Vacancies*

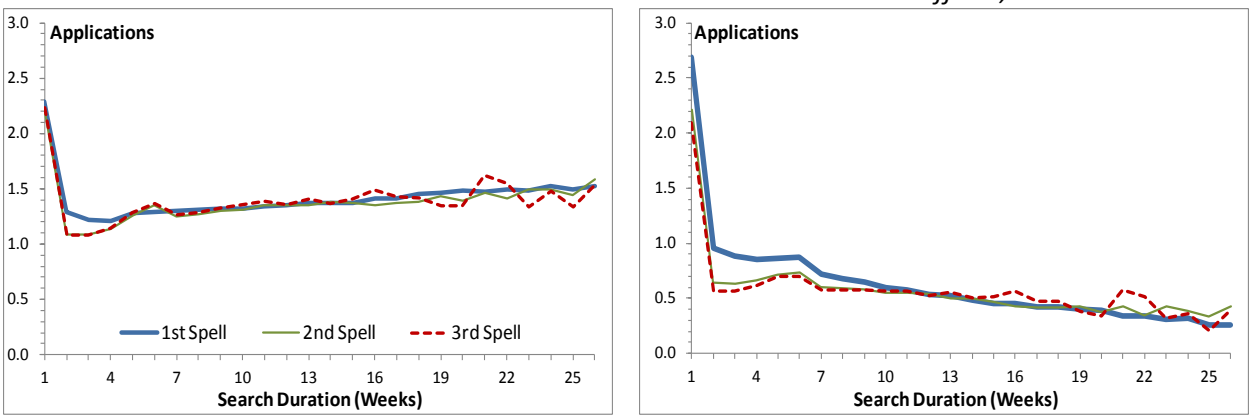

Notes: Figure shows the estimated relationship between applications per week and duration of search for our baseline model (left panel) and a model that additionally controls for active vacancies, fixed job seeker characteristics, and completed spell length (right panel). The model is estimated across all search spells for each job seeker.

Figure 12. Applications over the Duration of Search, Weak vs. Strong Local Labor Markets, Conditional on Demographics, Spell Length, and Local Vacancies

(a) High vs. Low Unemployment CBSAs

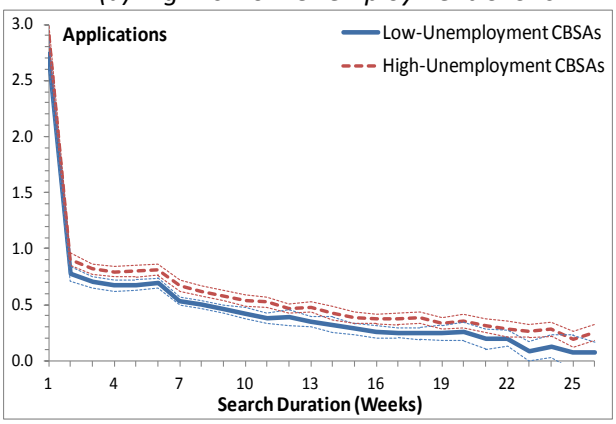

(b) High vs. Low Vacancy-Job seeker Ratio CBSAs

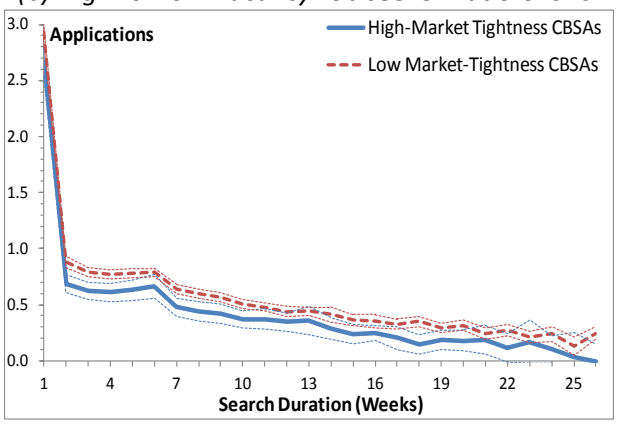

Notes: Figure shows the estimated relationship between applications per week and duration of search that controls for our baseline model's variables, as well as job seeker demographics, completed spell length, and the local number of active vacancies. The left panel reports the estimates for high vs. low unemployment metropolitan areas, defined as those in the top or bottom quartile of the CBSA unemployment distribution. The right panel reports the estimates for high vs. low market-tightness metropolitan areas, defined as those in the top or bottom quartile of the distribution ratio of vacancies to job seekers on the website. Dashed lines represent 95 percent confidence intervals (with standard errors clustered on job seeker's metropolitan area).

Figure 13. Search Behavior by Total Spell Length

(a) Mean Applications per Webpage View

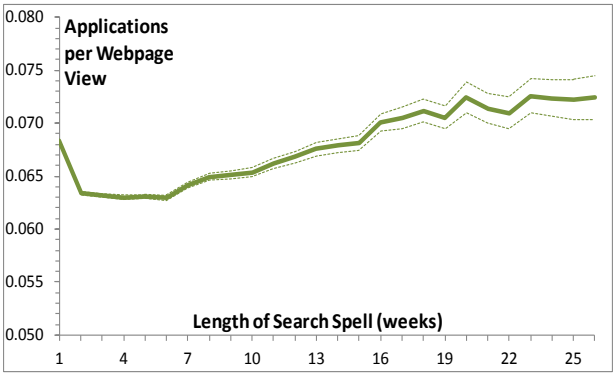

(b) Mean Time Spent Searching per Week

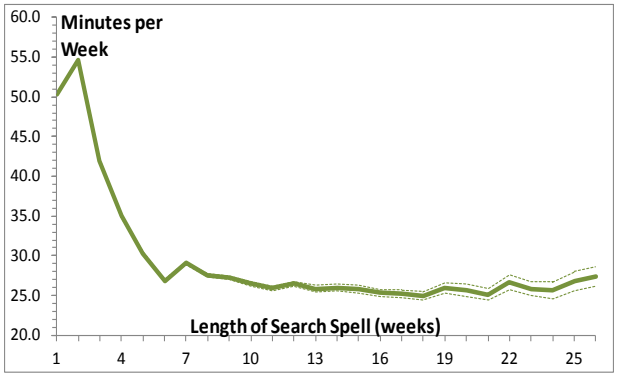

Notes: Panels depict the distribution of job seekers, the (unconditional) mean number of applications, the mean number of applications per webpage view, and the mean time spent searching the website (in minutes). Dashed lines represent 95 percent confidence intervals.

Appendix (not intended for publication)

Appendix A. Estimates without Spell Length Identification

Identifying when job seekers are actually searching as opposed to being idle on the website is an important part of our analysis. Many job seekers send applications over a short period of time, take a break from searching, and return later, sometimes many weeks later, to send applications again. We use a strategy of counting more than five weeks of inactivity as the end of one search spell and the start of a new spell. In the absence of this identification, results tend to be dominated by the fraction of job seekers who send zero applications, as the following figures show.

Figure A.1. Applications per Week, Counting Total Web Tenure as a Single Spell

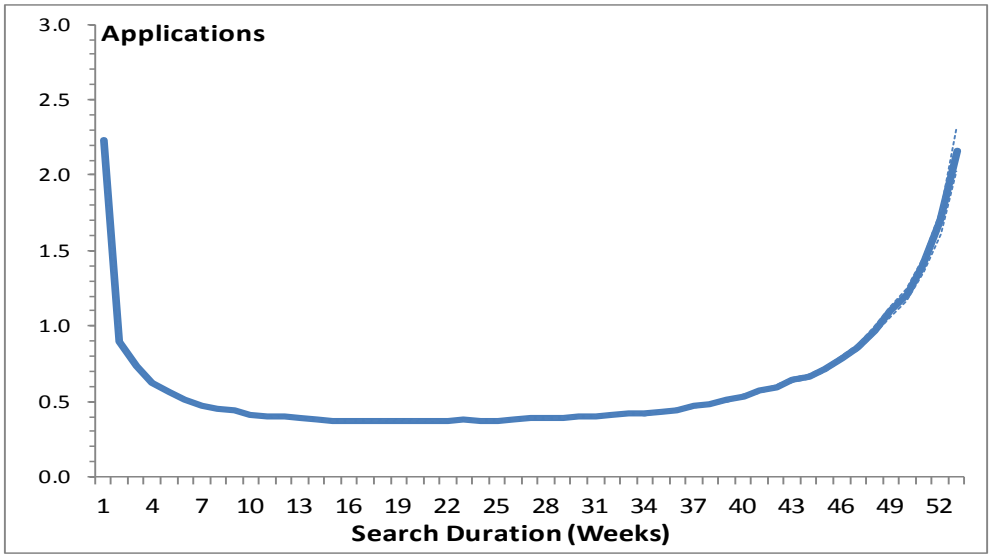

Figure A.2. Fraction of Observations with Zero Applications, Counting Total Web Tenure as a Single Spell

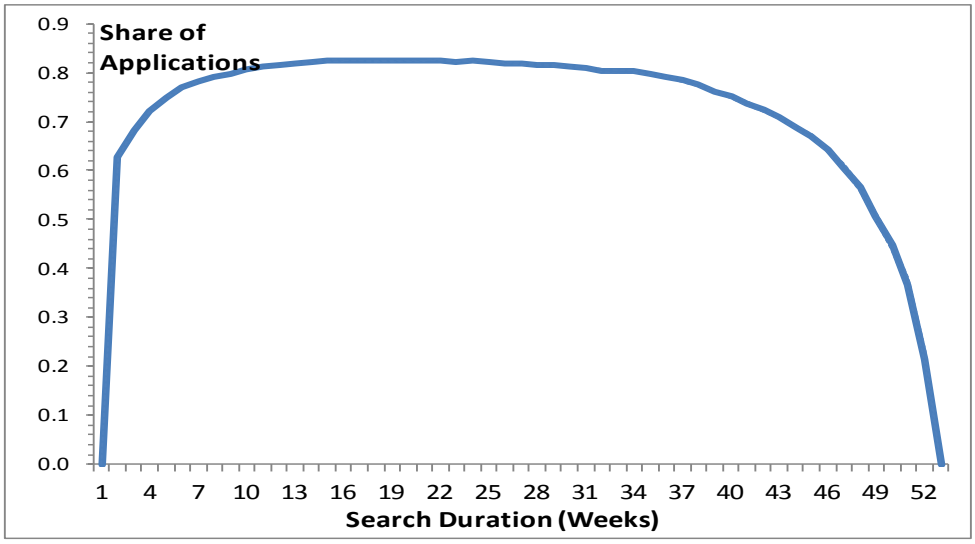

Figure A.3. Applications per Week by Total Web Tenure, Counting Total Web Tenure as a Single Spell

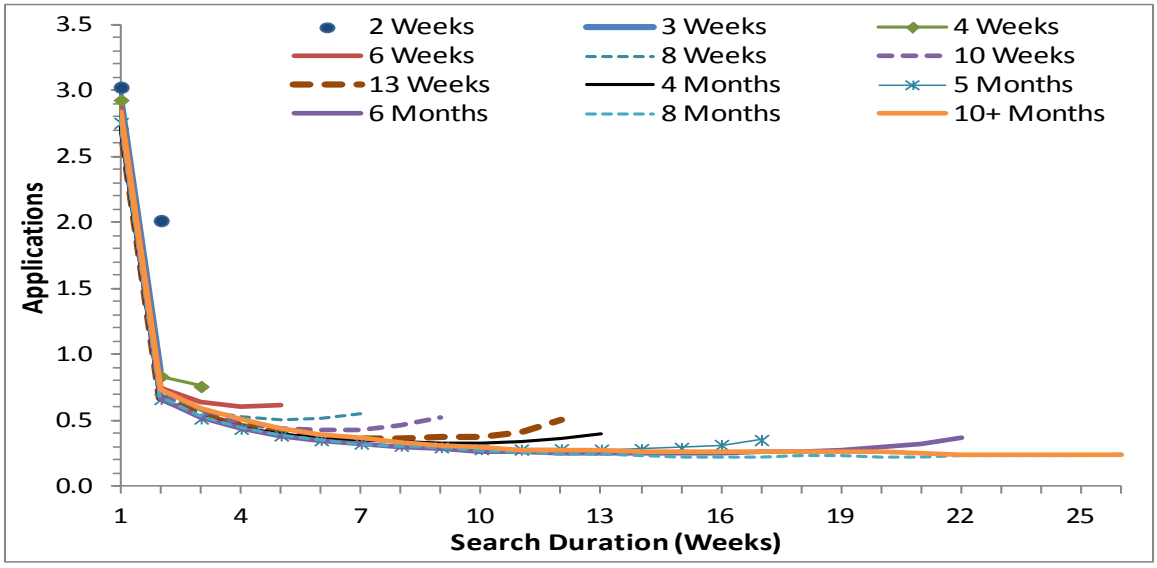

Appendix B. Estimates Using Differing Spell Length Criteria

We check the robustness of our cutoff of five weeks of inactivity by replicating our main results using a two-week and 13-week cutoff instead. The following figure replicates the baseline regression model of the analysis in the text, as well as the models that control for demographics and spell length, local vacancies, and the "full specification" of local vacancies and job seeker fixed effects. Each panel of the figure shows the estimated relationship between applications and search duration using the two-week, five-week, and 13-week cutoff rules.

The figures show clear quantitative differences in the estimated relationship between applications per week and search duration, especially in the specifications that do not include fixed effects for completed spell length or fixed job seeker characteristics. Qualitatively, however, the key result of our analysis holds regardless of spell length identification criteria. That is, applications decline with the length of search, especially after controlling for job seeker characteristics.

Figure B.1. Estimated Relation between Applications and Search Duration under Alternate Spell Length Identification Criteria

(a) Baseline Model

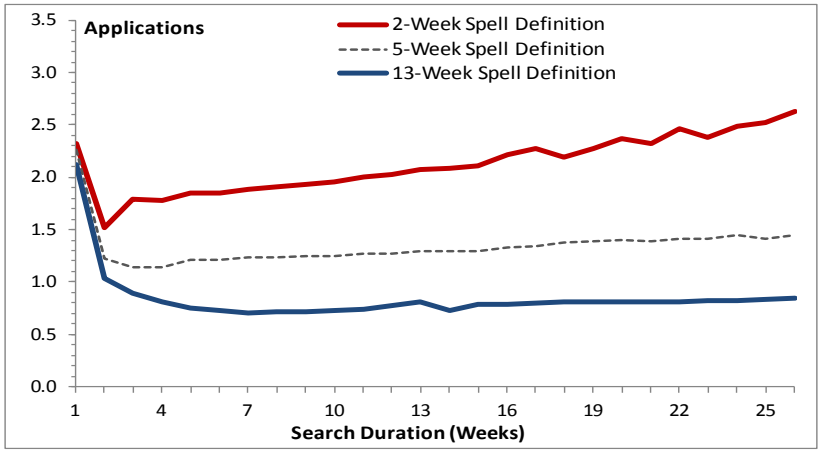

(b) Controlling for Demographics & Spell Length

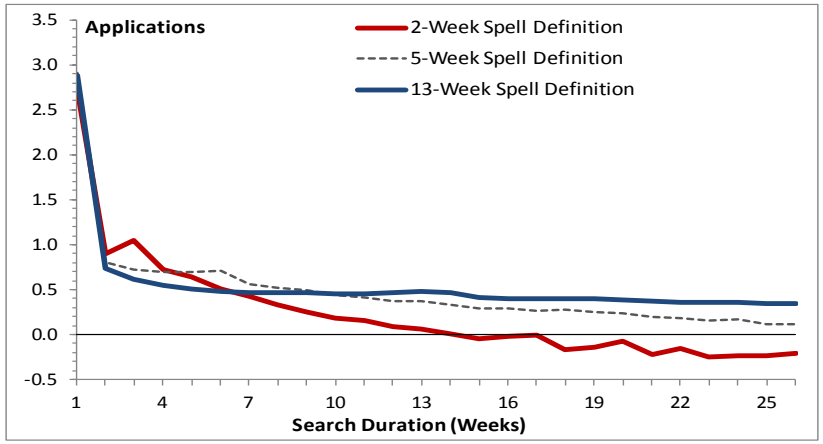

(c) Controlling for Local Vacancies

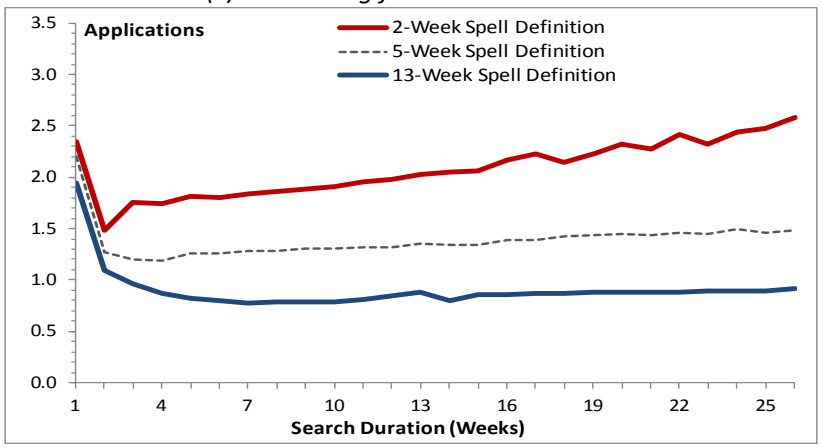

(d) Controlling for Local Vacancies & Job seeker Effects

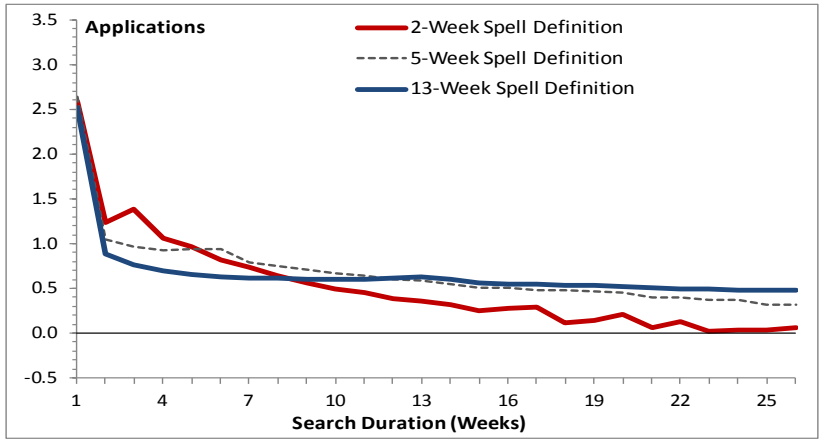

Notes: Panels depict the estimated relationship between applications sent per week and search duration under the four different regression specifications used in the analysis of the main text, using three different cutoff criteria to identify the end of a search spell: two weeks, five weeks, and 13 weeks of inactivity on the website.

Appendix C. Additional Results

C.1. Comparison to CPS Data

Table C.1. Differences in Duration, Website Sample and Current Population Survey, August 2011

	Duration of 1st Spell		Duration on Website		
	All Job seekers	*Job seekers with > 1 Application*	*All Job seekers*	*Job seekers with > 1 Application*	**CPS Unemployed**
Search or Unemployment Duration Distribution					
Less than 5 weeks	74.9	66.1	68.2	57.1	18.8
5-14 weeks	20.4	27.6	25.8	34.8	24.1
15-26 weeks	3.5	4.7	4.4	6.0	14.0
27 or more weeks	1.2	1.7	1.6	2.1	43.1
Mean duration, weeks	3.9	4.9	4.5	5.7	39.7
Median duration, weeks	1.0	3.0	1.0	3.0	20.6
N	187,354	138,626	187,354	138,626	*

Notes: Table reports the share of job seekers (or the unemployed, for the CPS) with an active search spell within the listed rage, with summary statistics on the duration of (incomplete) search spells included. Website data are from a cross-section of job seekers identified as actively searching during the CPS reference week of August 2011, and only include job seekers during their first identified search spell.
* CPS statistics are from published data, which typically come from a sample of about 100,000 individuals aged 16 and over.

Table C.1 compares the resulting distribution of search durations in our sample with the distribution of unemployment durations within the Current Population Survey (CPS). We use a cross section of job seekers during the CPS reference week of August 2011 for consistency with the CPS sample timing. As can be seen from the table, the average duration of the first search spell on the website is shorter than the duration of unemployment from the CPS. This occurs because the duration of the search on the website does not correspond to the notion of the duration of unemployment from the CPS. First, the job seekers in the sample include not only unemployed but also the employed and those who could have reported themselves as out of the labor force but still searched for work (e.g., retired individuals). Second, the unemployed job seekers might begin searching on the website a few weeks into their unemployment spell. Finally, the CPS unemployment duration measure faces issues with individuals transitioning between being unemployed and out of the labor force, i.e., unemployed

respondents may report their total time of non-employment as their unemployment duration, despite periods when search did not occur. Nevertheless, it is useful to understand how our measure of job seeker search spells compares with the search spells of the unemployed. From Table C.1, it is clear that the website has many more short-duration job seekers and much fewer long-duration job seekers relative to the unemployed in the CPS.

Figure C.1. Distribution of Job seekers by Completed Spell Length

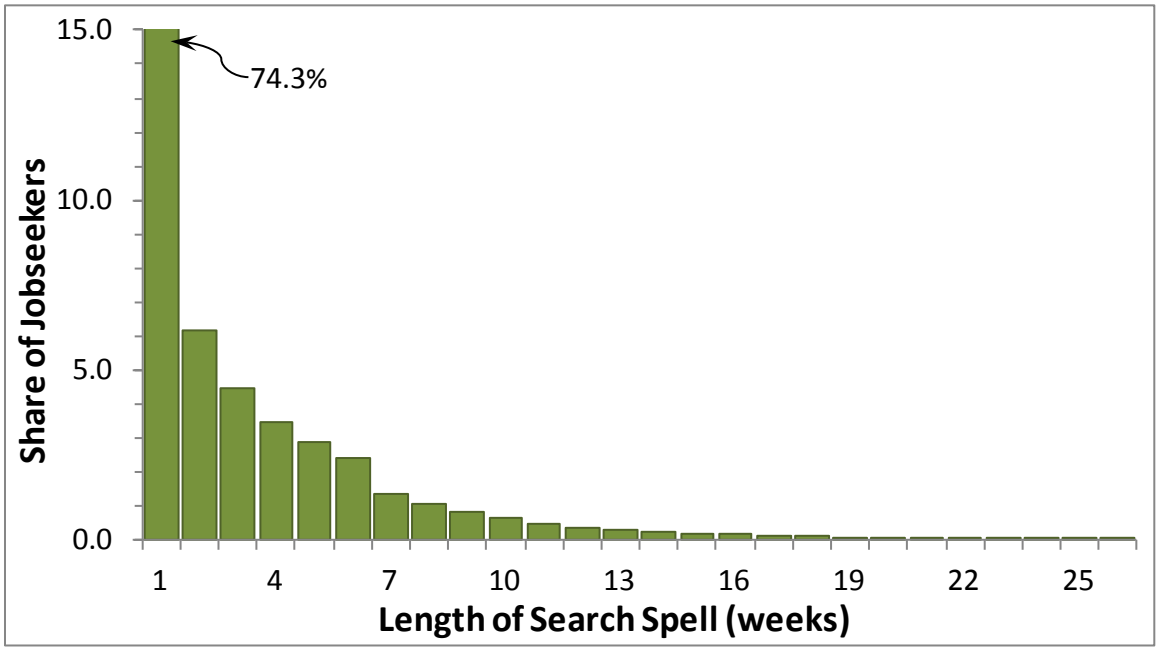

Figure C.2 shows the distribution of search duration for our sample of website job seekers, pooled across all periods. The distribution is nearly Pareto in shape, with most job seekers (74.3 percent) on the website for only one week.

C.2. Additional Results by Completed Spell Length

Figure C.2. Mean Applications per Week by Completed Spell Length

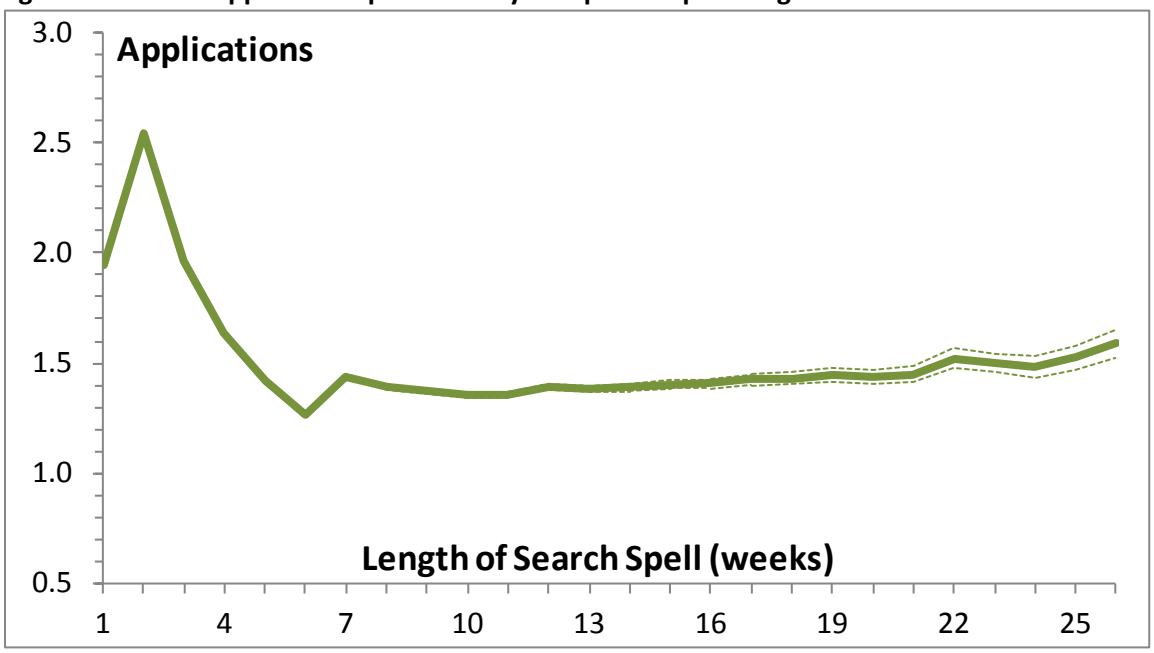

Figure C.2 shows that application behavior by completed spell length, before adding any controls, behaves similarly to application behavior by duration. Job seekers with the shortest durations send the most applications per week. Mean applications per week falls to its lowest level for those with spells of six weeks, then slowly rises thereafter. The similarity in pattern to the estimates in Figure 2 of the main text is not surprising, given that controlling for completed spell length has such a large effect on the results presented in Figure 3.

Figure C.2. Application Characteristics by Completed Spell Length

(a) Share to Modal Industry and Own Metropolitan Area

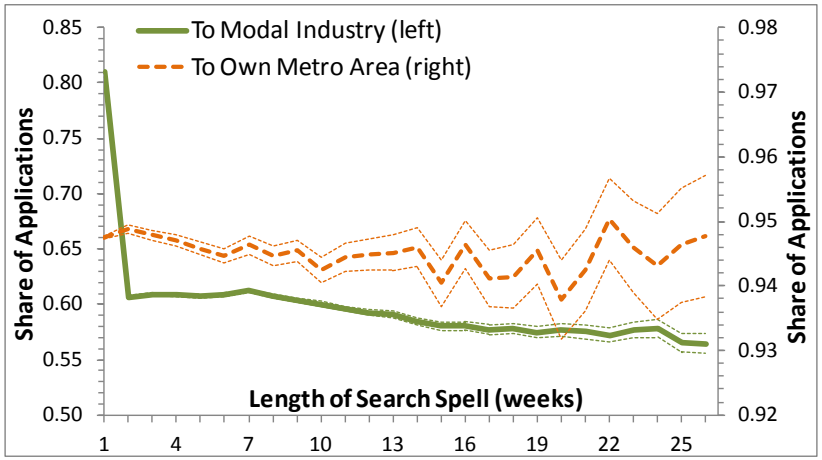

(b) Share to Previously Applied-To Vacancies

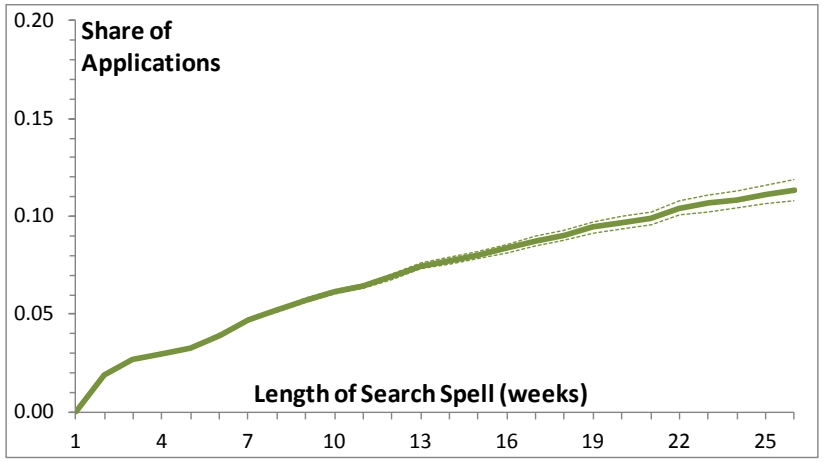

(c) Share to Newly Posted Vacancies

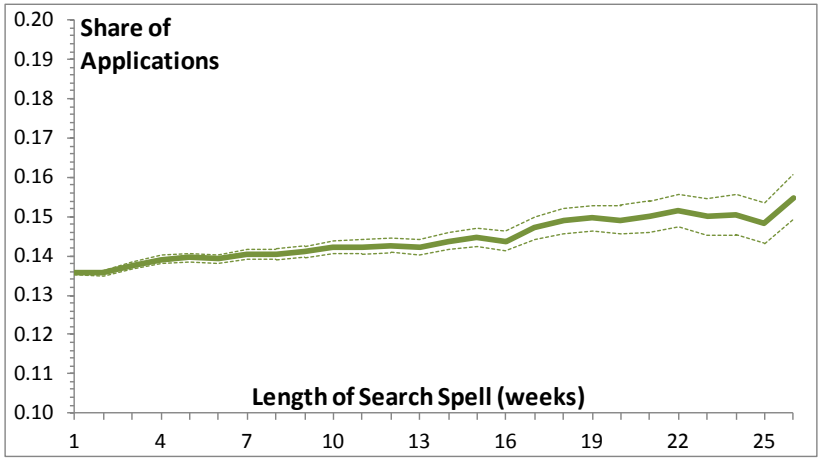

(d) Mean Duration and Total Applicants of Vacancies Applied To

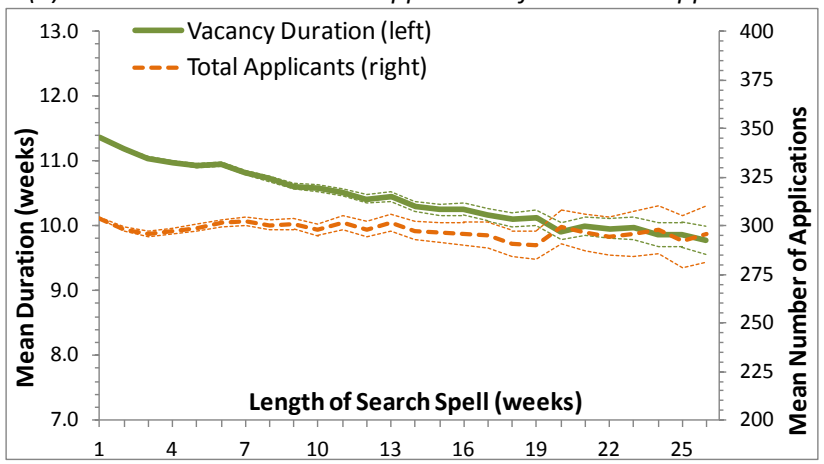

Notes: Panels depict the (unconditional) fraction of applications to a job seeker's modal industry, to the job seeker's metropolitan area, the mean number of applications per webpage view, and the mean time spent searching the website (in minutes). Dashed lines represent 95 percent confidence intervals.

C.3. Additional Results for Application Behavior, Conditional on Demographics, Spell Length, and Local Vacancies

Figure C.3. Application Characteristics by Duration, Additional Specifications

(a) Share to Modal Industry

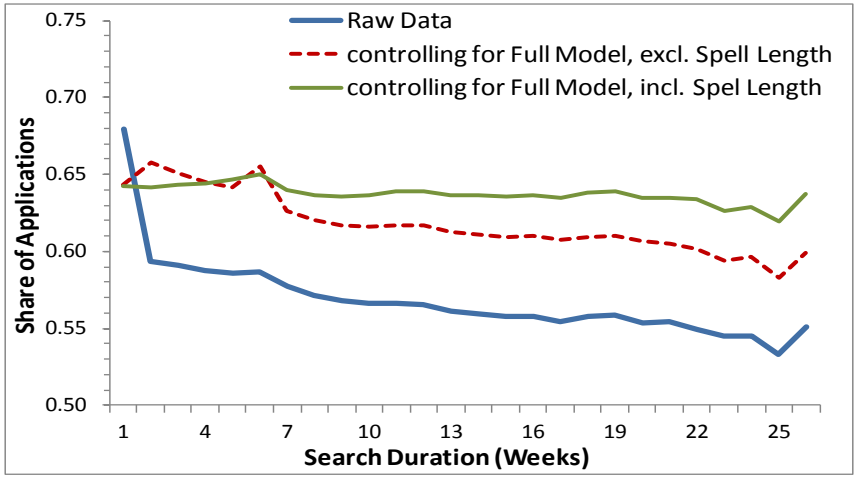

(b) Share to Own Metropolitan Area

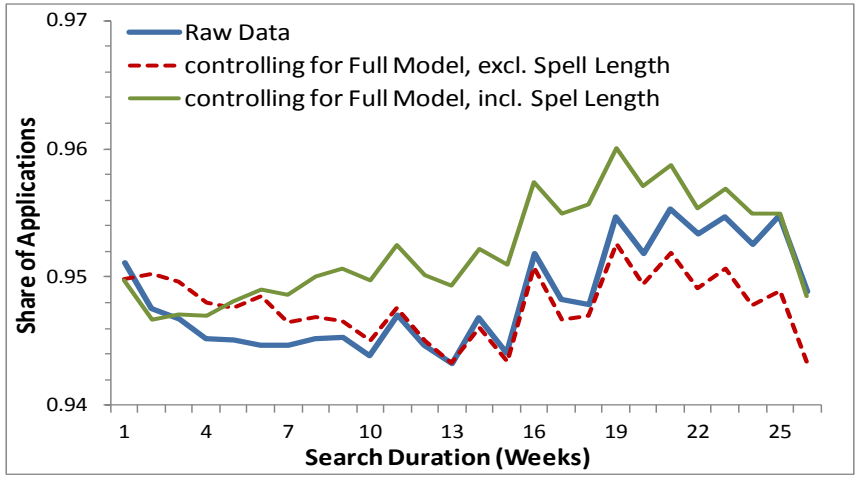

(c) Share to Newly Posted Vacancies

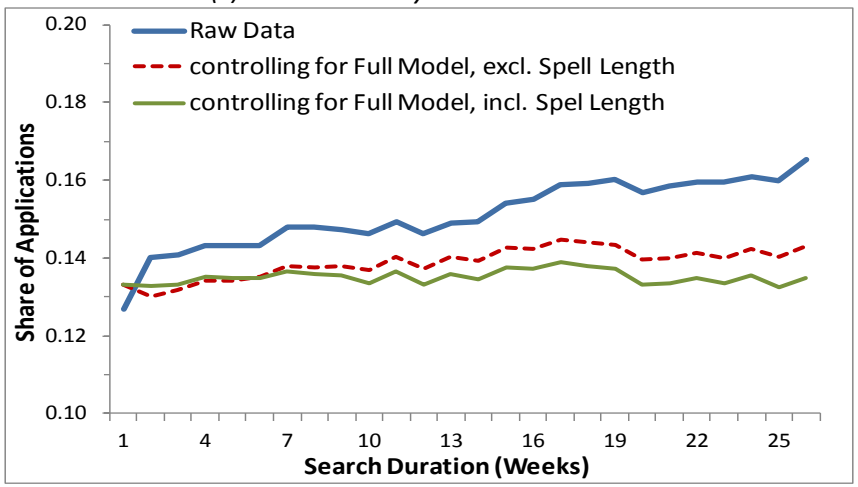

(d) Share to Previously Applied-To Vacancies

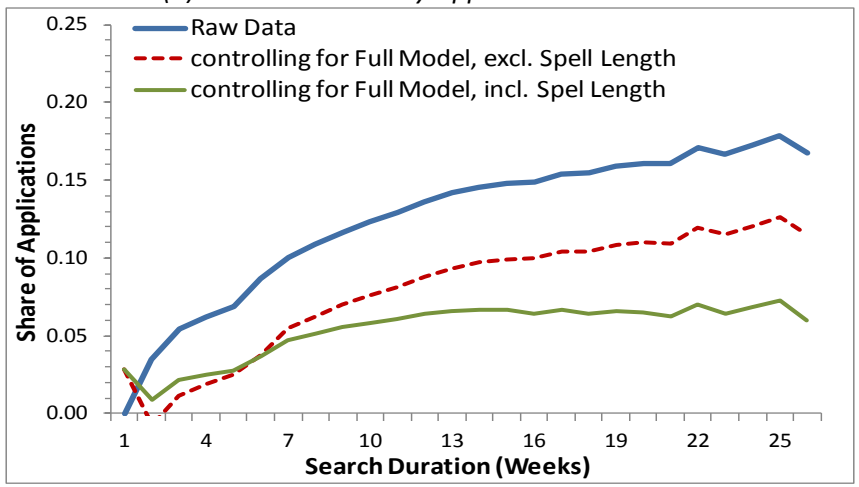

Notes: Panels depict the (unconditional) fraction of applications to a job seeker's modal industry, to the job seeker's metropolitan area, to newly-posted vacancies, and to previously applied-to vacancies by search duration, separately by the total spell length of job seekers.

Appendix D. Simulation Results, Constant Job Seeker Attrition Rate

This section presents additional results for our simulation of job seeker application behavior using the assumption that those who are marginally attached to search on the website quit using it entirely at a constant rather than declining rate over time. That is, we assume that $\rho(t) = \rho_0$, rather than $\rho(t) = \rho_0/t^{\rho_0+1}$, as is assumed in the main exercise. Since the new assumption is at odds with the sharply declining exit rates we see over the duration of search, we calibrate ρ_0 to the mean completed spell length of the full sample. Under the assumption of a constant quit rate, the inverse of average spell length equals $1-(1-f)^n+\ \theta\rho_0(1-f)^n$. All other calibrated parameters remain the same as in the text.

The results are in Figure D.1. As one can see, the results are qualitatively similar to those in Figure 9. The simulated full sample exhibits much greater dispersion across individuals of differing completed spell lengths in this case versus the case where the quit rate declines over time. This is because there are relatively more of the marginally attached remain on the website initially, are but they are also relatively more likely to quit the website later in their search spell. Despite this, the subsample of matches still shows considerably less dispersion across spell lengths. If our results were driven only by individuals dropping out of searching on the website, we should see less dispersion in the potential match sample when compared with the full sample. Our simulation under the assumption of a constant quit rate is also at odds with the data, reinforcing the implication that other forms of job seeker heterogeneity must drive the pattern of application behavior that we see in the data.

Figure D.1. Simulated Application Behavior by Completed Spell Length, Heterogeneous Tastes for Website Search

(a) Full Sample *(b) Simulated Match Sample*

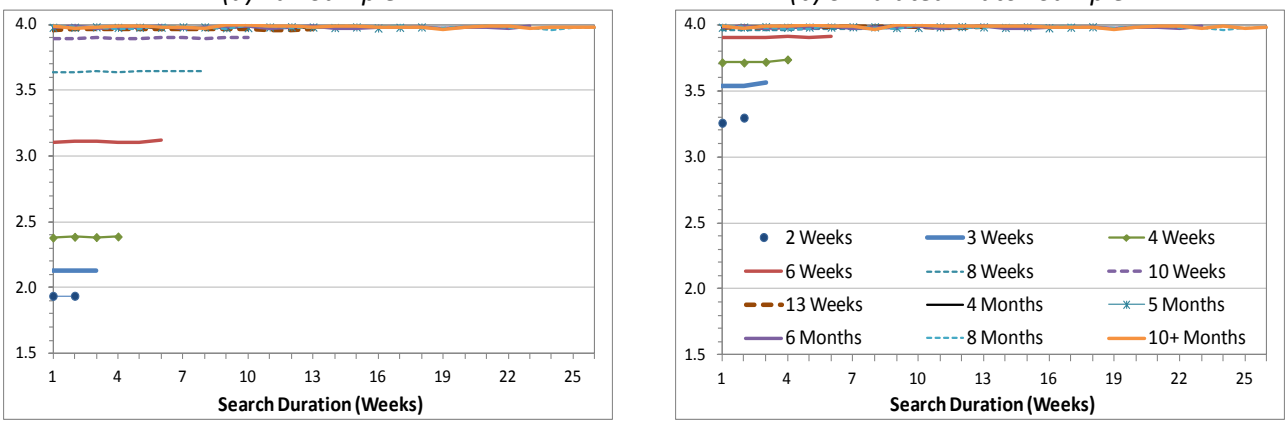

Notes: Figure shows the estimated (unconditional) relationship between applications per week and duration of search separately for job seekers based on the total length of their search spell using a simulated sample of job seekers calibrated to the empirical distribution of job seekers in our website sample. The left panel reports the estimates for all simulated job seekers, while the right panel reports the estimates for simulated job seekers who found employment through the website. Only selected spell lengths are reported.