Preliminary and Incomplete Draft, DO NOT CIRCULATE

10 July 2020

Application Flows

Steven J. Davis* and Brenda Samaniego de la Parra#

Abstract

We document new facts about recruiting and job search behavior in the US. We build a new U.S. database that links 125 million applications to over 7.5 million online job postings from January 2012 to December 2017. The raw data come from DHI Group, Inc., which owns and operates electronic platforms for posting vacancies and attracting applications. Postings fall mainly into computer-related occupations, technology sectors, financial services, and other occupations that require technical skills. We then assess the theoretical implications of our empirical results regarding firm and worker search. Our main findings are hard to reconcile with standard models of sequential employer search: First, the mean *posting* duration for single-position openings is 9.4 days, about one-fifth (23%) of the mean *vacancy* duration for comparable jobs in JOLTS data. Second, job seekers display a striking propensity to target new postings, directing almost half of applications to openings posted in the past 48 hours and more than three-fifths to those posted in the past 96 hours. Job-seekers concentrate their applications on their first day of search. Conditional on continuing to search, applicants submit decreasing batches of applications with an average 7-day waiting period. Labor market tightness, measured as mean applications per vacancy, does not have a large effect on vacancy posting duration. Platform functionality greatly affects the number of applications and their distribution.

Keywords:
JEL Codes:

Acknowledgments: We thank Yuri Bykov, Rachel Ceccarelli, Courtney Chamberlain, Jennifer Milan and Elizabeth Schillo for extensive consultations regarding the DHI data, online platforms, and business model. We thank Fernando Alvarez, Jason Faberman, Marianna Kudlyak, Ayşegül Şahin, Rob Shimer and seminar and conference participants at Arizona State University, Chicago, the 2018 IZA Conference on Matching Workers and Jobs Online in Bonn, the 2018 SAM Conference in Frankfurt, the San Francisco Fed, Stanford, UC-Berkeley, the Upjohn Institute, and Yale for many helpful comments on earlier drafts. The authors have received compensation from DHI Group, Inc. for developing and analyzing the DHI Database. Davis was previously a Senior Research Adviser to DHI.

* University of Chicago Booth School of Business and National Bureau of Economic Research. Steven.Davis@ChicagoBooth.edu.
# University of California at Santa Cruz, Department of Economics. bsamanie@ucsc.edu

I. Introduction *To be revised*

Application flows for job openings contain useful information about labor market conditions and the employer-worker matching process. Consider: First, vacancies attract larger application flows in slack labor markets and smaller flows under tight conditions.[1] This observation underscores the potential for using application flows to assess labor market tightness. Second, the timing and evolution of application flows for individual job openings influences employers' optimal search strategy. In particular, a heavy bunching of applications shortly after posting weighs in favor of a non-sequential search strategy. Third, short posting durations relative to vacancy durations indicate that the applicant gathering part of the search and matching process must be short compared to the parts devoted to screening, selection and post-offer recruiting. Fourth, from the applicant's perspective, if firms are searching non-sequentially, it makes sense for them to also implement a non-sequential search strategy and apply to several job simultaneously.[2] These observations indicate that application flows can be useful to evaluate the underlying assumptions of theoretical models of job search. Fifth, the search and matching functionality of online platforms affects both the volume of application flows and their distribution over vacancy postings and job seekers.

To pursue these ideas and observations, we first build a new U.S. database that links application flows to millions of online vacancy postings from January 2012 to December 2017. Our raw data come from the Dice.com platform of DHI Group, Inc., a company that owns and operates several online platforms for posting job vacancies and attracting applications. Employer-side clients of Dice.com comprise organizations that directly hire their own workers, recruitment firms that solicit applicants for third parties, and staffing firms that hire workers for lease to other firms. The DHI data identify employer-side clients, application flows for each vacancy posting, the job title, the city of employment, and other information about the job and each of its applicants. Vacancy

[1] This common-sense claim finds support in the strongly pro-cyclical behavior of vacancy durations, given that smaller application flows tend to produce longer vacancy durations. See Davis, Faberman and Haltiwanger (2012, 2013), Crane et al. (2016) and the monthly DHI Hiring Indicators report at http://dhihiringindicators.com for extensive evidence on the pro-cyclicality of vacancy durations and for comparisons to other tightness indicators.

[2] Van Ommeren and Russo (2009) discuss how firm's non-sequential search induces workers to also search non-sequentially. Burda and Profit (1996) and Lang and Majumbar (2004) discuss the implications of non-sequential search on the applicant side, including observing multiple simultaneous applications.

postings are concentrated in technology sectors, software development, other computer-related occupations, engineering, financial services, business and management consulting, and a variety of other occupations that require technical skills.

A few recent studies exploit data from online job boards to analyze various aspects of worker and employer search. Banfi and Villena-Roldan (2015) study the impact of explicit and tacit information about offer wages in vacancy postings on application flows in Chile. Marinescu and Rathelot (2015) use data on applications and vacancies to quantify the contribution of geographic mismatch to U.S. unemployment in 2012. Marinescu and Wolthoff (2015) study how applicant numbers and quality vary with job title and compensation information in online postings from early 2011 for Chicago and Washington, DC. Faberman and Kudlyak (2016) investigate how the intensity of online job search, as measured by an individual's application frequency, varies with the duration of job search.

We contribute to this literature exploiting the volume and granularity of the DHI data, coupled with its second-by-second tracking of postings and application flows. These characteristics of the DHI data allow us to construct novel indicators for a broad range of labor market outcomes across job categories[3] and firm types. This rich database yields new evidence on how the volume of application flows and the daily rate of applications per posting vary with posting age and search duration, day of the week, day of the month, employer type, and job type. Additionally, using a change in the platform's functionality, we are able to evaluate the effects of online job boards on the volume and distribution of application flows.

Previous empirical work weighed the evidence for different search strategies by employers and job seekers. Van Ours and Ridder (1992, 1993) provide early analyses of application flows and vacancy postings using Dutch data. Other early studies include Barron, Berger and Black (1997), Manning (2000), Weber (2000) and Russo, Hassink and Gorter (2005). In a more recent study, van Ours and Ridder (2009) point out that if employer search is sequential the number of rejected applicants is proportional to the number of vacancies. Therefore, testing the proportionality of rejected applicants to total vacancies is equivalent to testing for sequential search. They conclude

[3] Using text in the online postings, we group vacancies into close to 2,000 job titles, each with at least 100 distinct postings. For some purposes, we further group job openings into broader functional categories (e.g., software developer, project manager, business analyst) or skill categories (e.g., Javascript, Oracle, Linux).

that sequential search is rejected when firms use advertising or employment agencies to hire workers, but they cannot reject sequential search for other hiring methods. This result is consistent with the idea that if the time required to collect a set of suitable applicants is low (as is the case when employment agencies intervene), non-sequential search is optimal for employers.

Andrews et al. (2008) analyze the arrival rate of offers and job filling probability in the UK to assess the evidence for sequential search by employers. If search is non-sequential, the initial applicant arrival rate is high but the probability of filling the vacancy is low, since employers are sorting through all applications before making an offer. Since they find that the hazard rate of filling a vacancy is highest on the first vacancy day, they conclude sequential search cannot be ruled out. Also using data for the UK, Coles and Petrongolo (2008) find that the job finding rate for job seekers who have been unemployed for more than one month is mostly influenced by the inflow on new vacancies, rather than by the existing vacancy stock. They argue that this provides evidence in favor of stock-flow matching models.

We contribute to this literature by using the DHI data to link vacancy postings to the full set of applicant flows. The granularity of the data allows us to observe the exact duration of each posting as well as the exact arrival time of each application, down to the second. We can directly observe the behavior of application flows by vacancy age, as well as by applicants' search duration. Therefore, we can study whether the empirical evidence supports the predictions of leading theoretical models regarding the behavior of application flows and search strategies for employers and job seekers in the US.

After describing the DHI Database, we document several new findings about employer and worker search processes. First, "recruitment firms" (which solicit applicants for third parties) and "staffing firms" (which hire employees for lease to other firms) account for 67 percent of the vacancy postings in our data and attract 62 percent of the applications. This finding underscores the huge role played by matching and staffing intermediaries in contemporary U.S. labor markets, at least for the types of occupations and employers covered by the DHI Database.

Second, job seekers display a striking propensity to target new vacancy postings: 47 percent of applications flow to vacancies posted in the past 48 hours and 63 percent go to those posted in the past 96 hours. Applications per vacancy per unit time drop sharply as postings age.[4] Taken in

[4] This statement and the preceding one pertain to "standard" vacancy postings in the DHI Database, which account for 75 percent of all postings. As discussed in Section II.B, other postings reflect

isolation, these facts support the empirical relevance of stock-flow matching models, as in Coles and Smith (1998), Gregg and Petrongolo (2005), and Coles and Petrongolo (2008). However, the heavy bunching of application flows shortly after posting spells commence also weighs in favor of a non-sequential search strategy by employers, especially if coupled with a longer vacancy duration relative to posting duration. See Gal, Landsberger and Levykson (1981), Morgan (1983), Morgan and Manning (1985) and Andrews et. al (2008) on this point.

Third, *postings* for single-position openings are typically short-lived, with a mean duration of only 9.4 days. The mean *vacancy* duration for comparable jobs in the Job Openings and Labor Turnover Survey is more than four times longer. That is, the "search" phase of the hiring process, during which employers solicit and accept applications, is far shorter than the "selection" phase, which entails screening and interviewing applicants, selecting one for a job offer, extending an offer, negotiating terms, and waiting for a decision to accept or reject the offer. Many leading theoretical models of hiring behavior and frictional unemployment presume that employers follow sequential search strategies. This presumption is hard to square with the empirical pattern of a brief employer search phase and a much longer selection phase. We see our evidence on the brevity of posting durations and the heavy bunching of applications shortly after posting as strong motivation for renewed attention to search and matching models that feature non-sequential employer search.

Fourth, job seekers concentrate their application on their first day of search, submitting an average of 3.5 applications. Afterwards, mean applications drop sharply and they bunch their applications in weekly intervals. As noted by van Ommeren and Russo (2009), if firms search non-sequentially, applicants will find it optimal to also use a non-sequential strategy. This will result in sending multiple applications at once, waiting to hear back, and then applying again in batches if the initial search was not successful. Our empirical findings lend support to non-sequential search by both firms and job seekers.

We find that labor market tightness, measured as daily application flows controlling for realized application arrival, has a tiny, near zero effect on posting duration. Based on this finding, we argue that the procyclicality of vacancy duration likely reflects variation in the length of the selection phase of the hiring process, rather than in the meeting phase.

employers with recurring hiring needs for certain positions and recruiting firms that more or less continuously seek applicants for certain types of jobs. Each of these other postings typically involves multiple job openings, which greatly clouds the interpretation of posting age.

Exploiting changes DHI implemented to its platform on December 2014, we find that platform functionality greatly affects application volume and their distribution across jobs. Improvements in search technology and the ability to attract applicants to their postings disproportionally benefits smaller employers. The observed increase in mean daily applications per vacancy after the platform change decreases with firm size.

The paper proceeds as follows. Section II introduces the DHI Database, highlights key features that influence our measurement methods, and describes important changes in the functionality of DHI platforms during our sample period. Section III documents several basic facts about application flows and vacancy postings and considers how these facts relate to leading theories of labor market search, matching and hiring. Section IV considers the effects of changes in the search and matching functionality of DHI platforms on the volume and distribution of applications. Section V offers concluding remarks.

II. Dice.com and the DHI Vacancy and Application Flow Database

We build a new U.S. database that links 125 million applications to over 7.5 million online job postings and nearly 60,000 employer-side clients from January 2012 to December 2017. We worked closely with staff at DHI Group, Inc. to build and document the database, which draws on DHI's proprietary records for its Dice.com platform.[5] Before describing the database, we provide some useful background about the Dice.com revenue model, user experience and user profile.

1. The Dice.com Revenue and Pricing Model

Dice.com generates revenues from employer-side clients for vacancy postings and ancillary services, access to resume banks, and other recruitment services. Job seekers on Dice.com can register, review vacancy postings, and submit applications free of charge. They can also freely access Dice.com career development tools and content about local-market skill trends and salaries.

During our sample period, 98% of job vacancies on Dice.com were posted under "Subscription" contracts that grant the employer-side client a specified number of "job slots." This contract lets the client freely allocate different job postings to a given slot. However, the number of job postings visible to job seekers at a time cannot exceed the number of contracted job slots. The contract price varies with the number of slots and ancillary services. For example, DHI charges

[5] Other labor market matching platforms that DHI owned and operated during our period include eFinancialCareers, Biospace, Rigzone, ClearanceJobs.com, HealthECareers.com, and Hcareers.

extra to scrape job postings from the client's website and repost them on Dice.com. DHI offered other vacancy posting options during our sample period, but they accounted for tiny shares of all postings.[6]

Given the pricing of job slots, there is an implicit cost of keeping a given posting in active status, i.e., visible to job seekers. In particular, an active posting prevents the employer-side client from simultaneously using the job slot to post a different vacancy. Even when the cap on slots is nonbinding, the employer-side client has incentives to remove stale postings. For one thing, it is costly to respond to applicants. For another, the employer-side client opens itself to reputational damage when it leaves stale vacancies in active status. This reputational concern is important for employer-side clients according to DHI staff, partly because repeat contacts between particular job seekers and employer-side are common. In line with these remarks, we show below that posting durations on Dice.com are typically short, much shorter than vacancy durations for similar jobs. Thus, we think our measured posting durations closely approximate true posting durations for open job positions. In contrast, stale postings occur frequently on many prominent online platforms for posting job vacancies, and they create distinctive matching frictions and information externalities. See Cheron and Decreuse (2017) and Albrecht, Decreuse and Vroman (2019).

2. The Dice.com User Experience and User Profile

Dice.com visitors can browse and search postings by job title, location of employment, company name, skill requirements and other job characteristics. Browsing and searching do not require registration, but a visitor must register before applying for a job through the Dice.com platform. By registering, a Dice.com user can also post a resume and supply other information that is potentially useful in searching for jobs and in attracting interest from prospective employers. According to SEC filings, 81% of the job seekers who post resumes on Dice.com have a Bachelor's or more advanced degree. Over 70% have more than five years of experience, half have more than 10 years of experience, and most are employed (DHI Group, Inc., 2016, page 19).

DHI implemented several significant changes to its Dice.com platform in December 2014, with the goal of ensuring that "the right type of application is able to find and apply for the right

[6] Under its "Webstore" option, for example, an employer could purchase 1 to 10 credits. Each credit could be used to post a single vacancy for up to 30 days in the following 12 months. This option accounted for less than 1% of postings.

type of job."[7] First, it streamlined the registration process for Dice.com visitors, which led the number of registered users per unique visitor to rise by roughly one-third. Second, it removed information from job postings (e.g., contact information for hiring managers) that, in some cases, had facilitated applications outside of Dice.com. Third, the Dice.com platform upgraded to a new, more powerful search engine that enabled job seekers to better tailor their search queries and more easily identify jobs of interest. Fourth, registered users who complete an online profile can, since December 2014, make their information accessible to prospective employers. Employers can then signal interest to the job seeker, alerting him or her to a particular posting and encouraging an application. Finally, for job seekers who register and complete a profile, DHI streamlined the process of submitting certain applications. In many cases, a registered user with a completed profile can now submit applications in as little as ten seconds. For more information on the Dice.com user experience and the December 2014 changes in platform functionality, see Appendix B.

By making the Dice.com platform more convenient and useful for job seekers, we expect more applications to flow through the platform for any given number of postings, conditional on labor market tightness. Indeed, we find very large increases in the average number of applications per vacancy posting following these changes, as we discuss below.[8] We also explore whether and how enhancements to the search and matching functionality of the Dice.com platform altered the distribution of applications over vacancy postings, employers, and job seekers.

3. An Overview of the DHI Database

The DHI Database identifies employer-side clients and records when they post and withdraw particular vacancies. The database includes information about the client's industry, size, organization type, city and state. For each vacancy posting, we know the city of employment for the job on offer, the client's description of the job title as it appears in the online posting, a unique Job ID that links to the employer's unique Account ID, and the date-time stamp for each instance in which someone submits an online application, or clicks through to an external URL that accepts

[7] DHI implemented smaller changes to the functionality of its Dice.com platform at other times during our sample period. We obtained our understanding about the Dice.com changes through various conversations with DHI staff.

[8] The change in DHI platform prevents us from directly analyzing the evolution of labor market tightness since 2012. However, in a separate paper, we show how to extract useful information about the relative labor market slack by job title, job function, and skill requirements from microdata on vacancy postings and application flows.

applications for the posting. We also know the exact number of seconds a posting was active each day, and the number of views each posting received on each of these days.

While Dice.com serves employers and job seekers in many industries, the vacancy postings are concentrated in technology sectors, software development, other computer-related occupations, financial services, business and management consulting, engineering, and other technically-oriented professional occupations. This paper restricts attention to jobs in the United States, which account for 99% of the vacancy postings in the database.

"Direct hire" clients, which post vacancies to hire their own employees, make up 82 percent of the close to 60 thousand US-based client accounts in the DHI Database. Staffing and recruitment firms make up another 18 percent. Moreover, Direct Hire clients coupled with Recruitment and Staffing Firms post 99.7 percent of all vacancies and receive 99.8 percent of all applications. We drop the very small number of postings and applications associated with other types of client accounts. Staffing firms hire mainly with the aim of leasing employees to other firms. Recruitment firms seek suitable job applicants for their clients to consider hiring. Compared to other employer-side clients, recruitment firms are much more likely to use a single posting to recruit for multiple vacancies, jobs in more than one city, or even for multiple employers.[9]

When posting a vacancy, the DHI client chooses between two application channels. If the client selects the "email" channel, interested job seekers submit applications directly via the Dice.com platform. If the client chooses the "URL" channel, job seekers are redirected to an external URL operated by the client or a third party. The DHI Database records the number of completed applications via the email channel and the number of click-throughs to an external site for the URL applications. Thus, in the case of URL applications, we do not see whether the job seeker completes the application. The client can select different application channels for different postings. He can even switch application channel after posting, but that happens rarely.

As reported in Row (1) of Table 1, the DHI Database contains 7.5 million unique vacancy postings from January 2012 through December 2017. These postings attracted 125.5 million

[9] We developed these understandings through conversations with DHI managers and staff who work directly with DHI clients.

applications through DHI platforms during the same period.[10,11] Although recruitment and staffing firms make up less than one-fifth of DHI employer-side clients, they account for 67 percent of the unique vacancy postings and draw 60 percent of the applications. Email applications (i.e., those submitted directly through the DHI platform) account for 62 percent of all application flows, and URL applications account for the rest.

Many vacancies in the DHI Database have short offline spells, whereby a given Job ID (i.e. a distinct vacancy posting) is first posted, then taken offline and hence made invisible to job seekers for hours or days, and then made visible again. Short offline spells arise for various reasons: the client wants to check the content and appearance of a vacancy posting before starting to accept applications, the client briefly withdraws a posting to modify its description, or the client temporarily removes the posting as it screens a batch of applicants or awaits the outcome of an employment offer. For a given Job ID and date, we measure both elapsed calendar time since initial posting *and* cumulative time-to-date online net of offline spells. Both measures of posting duration are useful.

See Davis and Samaniego de la Parra (2019) for a full description of the DHI Database, including information about file structures, record layouts, variable definitions, and additional descriptive statistics.

4. Standard and "Long-Duration" Postings

Eighty percent of vacancy postings, individually identified through their Job ID's, exhibit the following pattern: (a) The client posts a vacancy on the DHI site, (b) a large majority of applications arrive within the first week or two after posting, and (c) the client permanently removes the vacancy posting within one month after first posting. The data exhibit variations on this standard pattern, but the key feature is the limited duration of the posting spell. For Job IDs that fit the standard pattern, we interpret each Job ID as a unique vacancy posting for a single opening. Other Job IDs do not conform to this pattern; instead, they remain online for many weeks or months, and

[10] About 0.2 percent of applications have a date-time stamp before the initial posting date of the corresponding vacancy or after its permanent withdrawal from the platform (i.e. its last active date). We drop these out-of-range applications on the view that they have an erroneous date-time stamp or an incorrect Job ID. Our results are robust to including them. Appendix E provides additional information on these out of range applications.

[11] We do not exclude duplicate applications, that is, applications submitted by the same applicant ID to the same job posting. We provide further information about these applications, and their impact on the distribution of applications across postings, in Appendix D.

applications flow in over time. Based on our examination of the data and our conversations with DHI staff, the vast majority of these "long-duration" postings reflect direct hire clients with ongoing hiring needs for certain jobs, and recruiting and staffing firms that continuously seek applicants for certain types of jobs.[12]

Given this understanding, we "slice" each long-duration posting into multiple postings, one for each calendar month during which the corresponding Job ID is active. This slicing operation lets us readily compare daily applications per vacancy, for example, within and between months. To operationalize this idea, let JobID_last denote the last date on which Job ID is active (i.e. the last day before the Job ID was permanently removed). If Job ID is **not** active on JobID_last-31, or at any earlier calendar date, we regard Job ID as a standard posting. If Job ID is active on JobID_last-31 or any earlier date, we interpret Job ID as a long-duration posting. In this case, we append a year-month identifier to Job ID to create a set of new unique Vacancy Posting IDs. For Job IDs that fit the standard pattern, we simply set Vacancy Posting ID=Job ID. Henceforth, we treat Vacancy Posting ID as our posting identifier, unless noted otherwise. As reported in Row (2) of Table 1, we have 11.7 million vacancy postings after this slicing operation.

The DHI Database records about 163 thousand vacancy postings and more than 1.75 million applications in an average month after slicing. Nearly half of the applications flow to standard postings, which account for 47 percent of vacancy postings after slicing. Some aspects of our analysis restrict attention to standard postings, because they typically pertain to a single job opening with a clearly defined first and last posting date.

5. The Distribution of Vacancy Postings by Employer Type and Size

Government employers and NGOs account for only 0.7 percent of direct hire postings in the DHI data. Accordingly, we interpret our results as pertaining to private sector hiring behavior. As reported in Table 2, direct hire postings are distributed widely by employer size. Over 90 percent of the direct hire postings are for jobs at privately held firms. In this regard, we note that privately held firms account for more than two-thirds of U.S. private sector employment (Davis, et al., 2007). In addition, because publicly listed firms are, on average, much larger and less volatile than privately held ones (Davis, et al., 2007), the share of vacancy postings and gross hires accounted for by listed firms is considerably smaller than its share of private sector employment.

[12] A small number of long-duration postings arise from single-position job vacancies that take many weeks or months to fill. This situation is rare for vacancies in our data, according to DHI staff.

6. Classifications by Job Title, Function and Skill Requirements

Marinescu and Wolthoff (2015) show the usefulness of job titles in classifying online vacancy postings. They find that job titles account for more than 90 percent of cross-vacancy variation in posted wages and more than 80 percent of the variation in the average experience and education of applicants. They also find that job titles are more useful in these respects than standard occupational classifications, because job titles contain more information about specialization, hierarchy (e.g., "staff accountant" versus "senior accountant"), and compensation.

We use job titles to classify vacancy postings and as controls in our statistical analysis. To do so, we first streamline and standardize the raw job descriptions as follows:[13]

1. We remove punctuation, information about location and text that aims to enliven the job or company. Examples of the latter include "*Unique Opportunity," "Amazing," "Exciting," "Innovative," "Start-up," "Urgent Need"* and similar terms.
2. We homogenize hierarchies into eight seniority indicators using text in the job postings: Level I, Level II, Level III, Level IV, Junior, Mid-level, Senior, Lead and Unspecified. We treat "*Entry Level"* as equivalent to Level I, "*Intermediate Level"* as equivalent to Mid-level, and *"Principal*" as equivalent to Lead.
3. We replace acronyms and standardize common terms. Examples include replacing "SDE" by "Software Development Engineer," "UAT" by "User Acceptance Tester," "DB Admin" by "Database Administrator," and "ROR" by "Ruby on Rails."

These steps yield more than 2 million job titles, most of which involve very few postings. We have 1,285 job titles with at least 250 distinct postings (Job IDs), 1,983 titles with at least 100 postings, and 2,746 with at least 50. As seen in Table 3, these frequently posted job titles account for over 90 percent of the Job IDs, Vacancy Posting IDs and applications in the database. Appendix Table A.1 lists the most common job titles in the DHI Database – such as Project Manager and Business Analyst among Direct Hire Clients and .Net Developer and Java Developer among Recruiting and Staffing Firms. Other common job titles include Software Engineer, Network Engineer, SAP Consultant, Systems Analyst, Program Manager, and Security Engineer.

For some purposes, we group postings and/or job titles into broader categories defined by Seniority, Job Function and/or Skill. "Seniority" refers to the hierarchy described in item (2) above.

[13] See Appendix A for a detailed description of how we streamline and standardize job descriptions.

Job Function refers to our grouping of (selected) job titles into 56 occupational categories such as "Programmer," "Developer," "Mechanical Engineer," "Electrical Engineer," "Consultant," "Analyst," "Business Analyst" and "Manager," among others.[14] Finally, our Skills classification reflects specific tasks or skills requirements, as specified in the text description. We consider 54 main skill categories such as "C," "SQL,", "Java," "User Interface" and "User Experience," and "Big Data."[15] In defining and selection Job Function and Skill categories, we prioritize job functions and skills that are most common in our database. When a job specifies more than one of the job functions or skills that we cover, we use the one that appears first in the text description, and we record the total number of distinct skills and job functions required by the vacancy posting.

III. Some Basic Facts and Their Implications

1. Vacancy Postings by Completed Spell Durations

Figure 1 shows the distributions of standard vacancy postings by completed spell duration, measured by time elapsed from initial posting date to final removal.[16] We compute duration in seconds and bin the results into 24-hour intervals, with Bin 1 corresponding to a spell duration of 24 hours or less. Pooling over observations for Direct Hire Clients and Recruitment & Staffing Firms, half of all standard postings last one week or less (summing the first 7 bins), and another 8 percent last more than 7 days but less than 8. Only 26 percent stay active for more than two weeks. The modal duration is 2 days (24-48 hours). The duration distribution for Recruitment & Staffing firm postings exhibits a pronounced second mode at 8 days (168 to 192 hours), while for Direct Hire clients the second most common duration is 1 days (up to 24 hours). The median posting duration is 7 days, and the mean is 9.4 days.[17]

[14] Appendix Table A2 lists the 56 Job Functions used to categorize vacancy postings.
[15] Appendix Table A3 lists the 54 Job Skills use to categorize vacancy postings. Appendix A also provides additional information on how we create the Job Function, Job Skill, and Job Seniority categories from the information included in vacancy postings' text descriptions.
[16] Some vacancies first appear online for less than 24 hours, draw no applications, and go offline for a spell before reposting. Based on discussion with DHI staff, we interpret these cases as trials that let the client inspect (and possibly modify) the posting before accepting applications. Accordingly, we exclude any initial spells that last less than 24 hours and receive no applications when calculating posting duration and age.
[17] Meaningful comparisons to online posting durations reported by other researchers are difficult due to measurement challenges and differences in pricing models across platforms. Consider two cases. First, Marinescu and Wolthoff (2015) consider point-in-time slices of CareerBuilder.com postings in early 2011. At that time, payment for a CareerBuilder.com posting covered a 30-day

The posting duration concept in Figure 1 differs from the JOLTS-based vacancy duration concept in Davis, Faberman and Haltiwanger (2012, 2013). DFH quantify the mean number of working days taken to fill vacant job positions, which involves much more than soliciting and accepting applications. It also involves screening and interviewing applicants, selecting an applicant for a job offer, extending an offer, negotiating terms, and waiting for a decision to accept or reject the offer. In this light, it is reassuring that the mean posting duration in the DHI Database is much shorter than the mean vacancy duration in JOLTS data. From January 2012 to December 2017, the mean JOLTS-based vacancy duration is 40.2 days for the Information sector, the closest JOLTS counterpart to jobs in the DHI Database.[18]

Table 4 provides information about how posting duration varies by job type and applications volume. Appendix Table C.1 provides analogous information with respect to company ownership and size.[19] Several results warrant attention. First, restricting attention to standard postings with job titles with at least 100 postings yields a nearly identical distribution of completed spell durations relative to considering all standard postings. Second, the median posting duration is a mere 7.0 days, and a quarter of all standard postings are active for 2.9 days or less. That is, the "meeting" phase of the search and matching process is very, very short for a larger share of postings. This characterization holds for all job types reported in Table 4, and it is broadly true of standard postings in the DHI Database. Third, and somewhat to our surprise, completed spell durations tend to rise with application numbers. We had anticipated that employers would lengthen posting durations in response to applicant scarcity. Of course, there is a mechanical effect cutting in the other direction: longer spells mean more time for application arrivals. We develop this point further

period. (These understandings reflect personal communications with Ioana Marinescu.) Marinescu and Wolthoff report a mean posting duration of 15.7 days, very close to the implied value if new postings arrive uniformly over the month and all postings remain listed for 30 days. Second, Brencic and Norris (2012) report a mean posting duration of 44 days in selected listings extracted from Monster.com in 2004 to 2006. During the period of their study, each payment for an online posting on Monster.com covered a 60-day period. They include postings that pertain to multiple job openings, which typically have much longer posting durations.

[18] We calculate this vacancy duration figure following the method developed by Steven J. Davis, R. Jason Faberman and John Haltiwanger (DFH) in "The Establishment-Level Behavior of Vacancies and Hiring" as described in the DHI Hiring Indicators reports at http://dhihiringindicators.com/. These reports contain mean vacancy duration statistics measured in working days, which we convert to calendar days by multiplying by (7/6).

[19] Appendix Figures D.5. and D.6 show the frequency distribution of posting duration for all postings (not just standard postings) by company ownership and size.

in section III.6, where we use application flows per vacancy to examine whether employers' lengthen posting duration as a response to applicant scarcity.

2. Application Flows by Posting Age, Application Volume, Job Type and Employer Size

Figure 2 displays the distribution of applications by posting age, defined as elapsed time since the posting first became active (i.e., visible to job seekers) to the time of application. As the figure shows, job seekers exhibit a striking propensity to direct applications to new and recently posted vacancies: 45 percent of applications flow to vacancies posted within the last 48 hours, and 60 percent go to those posted in the last 96 hours. Older postings attract relatively few applications.[20] Table 5 shows that a strong bunching of applications at freshly posted vacancies holds across quintiles defined by the volume of applications[21] and across a heterogeneous set of job categories. It is a ubiquitous feature of our data.

One reason fewer applications flow to older postings is because there are fewer of them. Recall that only 26 percent of standard postings in the database stay active for more than two weeks. To account for this fact, we also consider the relationship of application numbers to posting duration from a different angle. Specifically, Figure 3 shows how daily applications per vacancy posting vary with elapsed time since initial posting. Postings receive, on average, 2.1 applications in their first day online[22] and 2.4 applications on their second active day. Afterwards, applications per vacancy day drop sharply to 1.0 and even fewer per day as they age further.

Table 5 also reports equal-weighted (EW) and flow-weighted (FW) mean applications per vacancy, where the latter weights each posting by the number of total applications received. The equal-weighted mean reflects the central tendency of the application volume distribution from the employer perspective, while the flow-weighted mean reflects the central tendency from the applicant perspective. Table C.2 in the Appendix reports equal-weighted and flow-weighted medians.

[20] Appendix Figure C.3 displays very similar patterns when considering Direct Hires separately from Recruiting & Staffing Firms.

[21] To sort by applications volume, we first compute mean applications per vacancy at the job title level and then sort job titles into quintiles.

[22] It is important to note that during their first active day, a posting may (and often is) active for less than 24 hours. For example, a job posted on May 1st at 6:00 PM and removed on May 3rd at 10 AM. This posting is active for 6 hours during its first active day, 24 hours during its second active day, and for 10 hours on its third, and final, active day. Figure 3 shows mean applications per vacancy day (i.e. per active day) regardless of how many hours the posting was active for on each day.

For job titles with at least 100 postings, the mean number of applications per vacancy is 11.2 on an equal-weighted basis and 89.9 on a flow-weighted basis. Thus, the typical applicant faces many, many rivals for each sought-after job, even as employers face small applicant pools for most openings. Mechanically, as we discuss in further detail in the next section, this pattern reflects a highly uneven distribution of applications over postings (Figure 4). In terms of economics, this pattern is consistent with at least two somewhat different interpretations: first, that a modest share of vacancies is highly attractive to many job seekers because of high compensation, good working conditions, high job security or other desirable attributes. Second, that skill mismatch is an important phenomenon that curtails the size of applicant pools for many vacancies. Table 5 provides some impressionistic support for the latter interpretation in the small applicant numbers for Electrical Engineers and Mechanical Engineers, two occupations with demanding skill requirements. Meanwhile, the lower application volumes for postings for jobs in Sales and Business Development may reflect the former interpretation: within applicants in the DHI pool, these jobs are relatively less attractive. We explore the determinants of the heterogeneity in application flows across postings in more detail below.

Figure 5 shows weighted and unweighted mean applications per posting by employer size class. The figure restricts attention to direct hires, because employer size is less meaningful for recruitment and staffing firms. Perhaps surprisingly, there is no strong, simple relationship between employer size and the size of applicant pools. The largest employers draw the smallest applicant pools, while employers with 1,000 to 2,500 employees draw relatively large pools, especially on an unweighted basis. Direct hire clients with five to nine employees draw the highest applications per vacancy. Clients with zero reported employees draw relatively small applicant pools. These clients are likely a mix of shell companies and start-up firms. Controlling for differences in the mix of job titles (Panel C in Figure 5) does not greatly alter the relationship between employer size and mean applications per posting.

From the applicant perspective (Panel D in Figure 5), competition is similar at firms with 10 to 19 employees, 100 to 249 employees, and those with 1,000 to 2,499 employees, where the average applicant competes with an additional 8 job seekers. The average applicant for jobs at firms with 5 thousand employees or more competes with 17 fewer job seekers. Figure C.2 in the Appendix reports equal and flow weighted median applications per vacancy by employer size.

3. Highly Uneven Distribution of Application Flows

Figure 4 displays the distribution of standard postings by number of applications received in the first 14 days online (measured in 24-hour intervals). For Direct Hire clients, 19 percent of standard postings attract no applicants in the first 14 days and 13 percent draw only one. For Recruitment & Staffing firms, 23 percent attract no applications in the first 14 days and 15 percent draw just one. 14 percent of postings by Direct Hire clients and 10 percent of those by Recruitment & Staffing firms attract 20 or more applicants in the first 14 days. Figure C.1 in the Appendix tells a similar story for long-duration postings, echoing the message of Figure 4 for standard postings: Applications are distributed highly unevenly across vacancy postings.

From an economic perspective, it might seem surprising that many postings in our database draw few applications. Three observations are useful in this regard. First, most Dice.com postings have demanding technical qualifications such as Java developer, software engineer, systems administrator, SAP consultant, LINUX administrator, data scientist and electrical engineer. Second, the job postings are concentrated in occupations with relatively rapid demand growth in recent years, potentially outstripping the pace of skill adjustment on the supply side of the labor market. For both reasons, skill scarcities are more prevalent for the jobs on Dice.com than for jobs in the economy as a whole. Third, DHI takes steps to block visits from certain foreign locations and from IP addresses and User IDs with a history of nuisance applications. These steps are part of DHI's efforts to provide high-quality applicant pools to its employer-side clients.[23]

The share of postings with no applications, as well as the unevenness in application flows conditional on receiving any applications, are much too great to be explained by a model of random assignment. To see this point, consider the classic balls-into-bins problem, where we treat applications as balls and postings as bins. Suppose a applications flow randomly to v postings. The simple mean number of applications per posting is a/v. Assuming search is random, the number of applications received by a given posting follows a binomial distribution with parameters a and $(1/v)$. The expected fraction of vacancy postings that receive exactly x applications is given by $\frac{a!}{(a-x)!\,x!}\left(\frac{1}{v}\right)^{x}\left(1-\frac{1}{v}\right)^{a-x}$. The flow-weighted mean of applications per posting, weighting each

[23] DHI offers various packages to help clients identify and possibly filter "low quality applicants." During our sample period, most clients used the baseline package which involved no intervention by DHI beyond blocking suspicious foreign applications and identifying third-party applicants.

posting by its applications count, is $(a/v) + 1 - (1/v)$. Thus, the gap between the flow-weighted mean and one plus the simple mean is a measure of distance from random assignment.[24]

Table 6, column (1) shows moments from the distribution of applications per vacancy for Direct Hire Standard Postings from the DHI data. Column (2) shows the moments implied by a random assignment model (i.e. a binomial distribution) that matches the simple mean applications per vacancy observed in the data. Comparing these two columns evidences the fact that the dispersion in application flows and the share of postings that receive zero applications exhibited in the DHI data cannot be reconciled with a model in which applications are allocated randomly across all vacancy postings. A binomial distribution fixes the ratio of the median to mean to be equal to 1. Meanwhile, the data shows a median to simple mean ratio of 0.34. The variance in the distribution of applications per vacancy in the random assignment model, $\frac{a}{v}\left(1 - \frac{1}{v}\right)$, is a linear function of the mean and always smaller than the mean. As $\frac{1}{v} \to 0$, the ratio of the variance to the mean number of applications tends towards 1. Instead, the DHI data presents a much larger standard deviation.

To allow for additional flexibility in the relationship between the simple mean and the variance in application flows, we next consider a negative binomial (NIB) distribution. The NIB model directly allows for overdispersion, that is, for wider discrepancies from mean application flows between the data and those implied by a standard binomial distribution. Under a NIB distribution with mean a/v and dispersion parameter θ, mean applications per posting are a/v and the variance equals $a/v\big(1 + \theta(a/v)\big)$. The dispersion parameter, θ, scales the variance and makes it a quadratic rather than a linear function of the mean. Notice that when θ is equal to zero, the distribution collapses to a binomial distribution and, hence, the random assignment model is a special case of the negative binomial distribution.

Column (3) in table 6 shows the moments implied by a NIB model that targets the simple mean and the standard deviation of the empirical distribution of application flows.[25] By matching the standard deviation in the data, this model is able to perfectly match the ratio of the flow-weighted to simple mean. However, it predicts 2.5 times the share of postings with no applications. This larger than expected share of postings with zero applications is caused by two facts: first, the

[24] See Appendix F for a proof of this conclusion.
[25] The mean (μ) and overdispersion (θ) parameters of the NIB model that match the simple mean and standard deviation in the data are, respectively, 11.8 and 5.8.

standard deviation in the DHI data is more than twice the simple mean and, second, application count is restricted to be non-negative. Hence, when we target these two moments, a large share of postings are predicted to have zero applications.

The NIB model produces a much greater than observed incidence of zero-application postings relative to the empirical distribution. A model with greater flexibility that allows for two separate processes to discipline the overdispersion in application counts and the share of postings with zero application can potentially achieve a better fit of the data. We therefore next consider the moments implied by a zero-inflated negative binomial model (ZINB).

The ZINB model allows for two groups of postings: one that will never receive applications ("always zero" group) and a second one for which applications are distributed based on a negative binomial distribution, where zeros sometimes occur. Let π be the share of postings in the "always zero" group. The expected count of applications per vacancy, conditional on not being part of the "always zero" group, is $\frac{a}{v(1-\pi)}$. The variance in applications is given by $\frac{a}{v}(1-\pi)\left(1+\frac{a}{v}(\pi+\theta)\right)$. The share of postings with zero applications is $\pi+(1-\pi)\left(1+\theta\left(\frac{a}{v}\right)\right)^{-1/\theta}$. The ZINB distribution allows us to separately consider deviations from the random search model due to the higher than expected share of postings with no applications, and larger than expected variance in application count (beyond the concentration in zero application postings). It is also important to note that if none of the postings are of the "always zero applications" group, (i.e. if $\pi = 0$), the model collapses to a negative binomial.

Column (4) in table 6 reports the moments for the distribution of applications per vacancy posting for the ZINB model that best jointly approximates the simple mean, standard deviation, and share of postings with zero applications in the DHI data. We estimate the parameters of the model by minimizing the absolute value of the sum of the deviation between the data and the model's predictions for these 3 moments. The value for the overdispersion parameter, θ, necessary to predict the standard deviation in applications observed in the DHI data is less than half (5.8 vs. 2.4) once we introduce the zero-inflation parameter and we allow the simple mean to depart from its empirical counterpart.

Figure 6 shows the implied simple mean applications per vacancy, share of postings with no applications and standard deviation in applications by a zero-inflated negative binomial model as the overdispersion parameter increases. For each value of the overdispersion parameter, θ, we

choose the other two parameters of the ZINB model (the mean of the negative binomial distribution, μ, and the zero-inflated probability, π) to minimize the sum of the absolute value of the deviation between the model implied and the empirical simple mean and the absolute value of the deviation between the model implied and the empirical share of postings with zero applications. A random assignment model would predict a constant, $\left(1-\frac{1}{v}\right)$ ratio between the variance and the simple mean of application flows. Instead, the overdispersion parameter introduces flexibility to the dispersion in applications. Consistent with the results presented in table 6, a model with an overdispersion parameter equal to 2.24 matches the standard deviation (26.8) and the share of postings with zero applications (18.9%) observed in the data but overshoots the mean applications per vacancy (18.8 vs. 11.8 in the data).

In short, the data exhibit enormous departures from random assignment. Instead, models that allow for overdispersion in the distribution of applications are better able to match the DHI data. A NIB model perfectly matched the simple mean and the standard deviation in the data and its related moments (such as the ratio of flow-weighted to simple mean). A ZINB model is better able to target the data's dispersion and share of postings with no applications.

There is much heterogeneity across job postings that could explain the observed departure from random assignment in application flows. It is possible that although in the aggregate the distribution of applications does not seem to conform with random search, within more narrowly defined groups of postings (grouped by job or employer characteristics), application flows better conform with random assignment. The gap between flow and equal-weighted application flows we presented in table 5 for various job categories provides initial evidence contradicting this premise. We examine this in more detail next by estimating the residual overdispersion in the DHI data after controlling for various sources of heterogeneity across job postings. Estimated values of overdispersion close to zero, after controlling for relevant job and employer characteristics, would indicate that applications are randomly distributed within sets of postings that are observationally similar. Our findings suggest that non-random application distribution persists even after controlling for a wide range of heterogeneity.

Let A_j be the total applications received by posting j. Equation (1) shows the probability distribution for A_j in our baseline specification. For all specifications, we model the incidence of zeros as a logistic link function of an indicator variable equal to one for postings that have at least one view by job seekers. Job seekers must select (i.e. view) a posting in order to apply to it.

Therefore, receiving zero views is a good predictor of the "always zero applications" group. We separately model the distribution of application flows across postings, conditional on not being in the "always zero" group, using a negative binomial distribution.

$$\Pr(A_j = a) = \begin{cases} \pi_j + (1-\pi_j)g(A_j = 0) & if\ \ a = 0 \\ (1-\pi_j)g(A_j) & if\ a > 0 \end{cases} \quad (1)$$

where $\pi_j = \frac{\exp(\beta_0+\beta_1 I[views_j>0])+\epsilon_j}{1+\exp(\beta_0+\beta_1 I[views_j>0])+\epsilon_j}$,

$$g(A_j) = Pr(A_j = a \mid \mu_j, \theta)\frac{\Gamma(A_j+\theta^{-1})}{\Gamma(\theta)\Gamma(A_j+1)}\left(\frac{1}{1+\theta\mu_j}\right)^{\theta^{-1}}\left(\frac{\theta\mu_j}{1+\theta\mu_j}\right)^{A_j}$$

and $\mu_j = \exp(\alpha_0 + \alpha_1 X_{1,j} + \alpha_2 X_{2,j} + \cdots + \alpha_n X_{n,j} + \epsilon_j\})$

Table 7 shows the estimated coefficients for the negative binomial, α_j, progressively controlling for additional vacancy posting and employer characteristics that job seekers potentially use if directing their search. The baseline specification, presented in the first row, controls for a set of 55 skill requirements fixed effects, whether a wage is posted, whether the vacancy allows applications from 3rd parties,[26] whether the employer is a Direct Hire client (versus a Recruitment or Staffing Firm), and whether the posting receives applications directly through the Dice.com portal (email applications) or redirects the applicant to an external URL.[27] Each row adds the fixed effects listed in the first column as additional controls. We estimate the remaining overdispersion, θ, using maximum likelihood.

The extent to which the estimated overdispersion parameter changes as we control for job and employer characteristics is useful to identify the job and employer characteristics that job seekers target to direct their search. The estimated overdispersion for the last row, 1.33, indicates that even after removing the effect of having an explicit wage offer in the vacancy posting, variation across calendar time, employer size, state, narrowly defined job titles, and even employer fixed effects, the data exhibits an important departure from random assignment. As reference for comparison, at an overdispersion value of 1.33, the dispersion in application flows is 17 times

[26] We assume that job postings that do not explicitly prohibit applications from third parties allow them.

[27] We cannot directly observe whether a job posting allows for email or URL applications. Instead, this classification is based on the type of applications a posting receives. Therefore, we cannot determine the application channel for postings that do not receive any applications. Instead, we randomly assign job postings with zero applications into one of the two channels.

higher than the variance generated by a standard binomial distribution. We interpret the remaining overdispersion parameter as a measure of remaining directedness in search, conditional on observables.

Table 7 also provides information about the determinants of application flows. First, job postings that include an explicit numeric value (or range) for compensation do not attract an economically meaningful higher number of applicants, after controlling for skill requirements.[28] Second, jobs posted by Direct Hire clients receive 1.35 times ($e^{0.3}$) more applicants than Recruitment and Staffing Firms. Third, the magnitude of the effects of baseline job posting characteristics (wage posting, permission for 3rd party applicants, client type, and application channel) are relatively unchanged as we progressively add controls. Finally

We allow for heterogeneous departures from random search across skill categories by separately estimating the remaining overdispersion from a ZINB model for each skill category in table 8. Each row considers postings with the listed skill requirement. We progressively add the column title as an additional control in the negative binomial model. In an intercept only model (i.e. a model that only accounts for the inflated zero application count generated from vacancy postings that are not viewed by any job seekers), the estimates for the overdispersion parameter range from 1.4 to 2.9. These overdispersion values indicate a variance in application flows across postings with the same skill requirement ranging from 10 to 62 times the variance generated in a random assignment model.

The magnitude of the decrease in the overdispersion parameter as we control for additional vacancy posting and employer characteristics is indicative of the importance of the added control in job seekers' directed search. Market tightness, measured as month-year fixed effects within each skill category, are an important determinant of application flows across postings. After controlling for tightness, whether the vacancy posting allows for 3rd party applicants and employer fixed effects are the factors that most contribute to reducing overdispersion. However, even after adding the full set of controls and fixed effects, we reject the null hypothesis that the dispersion parameter equals zero, that is, that applications are randomly allocated after controlling for all observable characteristics,[29] at any meaningful significance level (p-value<0.0001).

[28] Application flows do not significantly change with the natural log of the posted real wage either.
[29] If the overdispersion parameter is equal to zero, the distribution collapses to a zero-inflated binomial distribution. In other words, conditional on not being part of the "always zero

4. *Posting Duration and Labor Market Tightness*

Previous research establishes the strong pro-cyclical behavior of vacancy durations.[30] This pro-cyclicality can be due to employers extending posting durations during booms when there are fewer available applicants. However, the evidence regarding the variation in vacancy duration across the cycle does not provide information about whether posting duration is also pro-cyclical. Even if posting duration were irresponsive to labor market tightness, or countercyclical, vacancy duration could increase during booms if employers become more demanding about whom they hire lengthening the screening and selection process.

First, we assess whether employers shorten posting duration upon higher realizations in daily application flows for each skill requirement category. Equation (2) shows our specification, where *j* denotes each vacancy posting, *s* denotes skill categories, *t* denotes the calendar time (month-year) when the posting was first active, and *f* denotes job functions. $I[Applications_j>0]$ is an indicator function equal to 1 if posting *j* received any applications. The dependent variable is the natural log of posting duration, measured time elapsed from the first to the last active date-time, expressed in 24-hour intervals. For the realization of daily applications, we use the total number of applications received in the first 14 postings days divided by the minimum of total posting duration and 15. We focus on standard postings with first active dates prior to December 1st, 2017 (to keep standard postings with completed spells only). We exclude postings in skill requirement categories with fewer than 25 distinct job postings in any given month and all postings

$$\ln\left(duration_j\right) = \beta_0 + \sum_s \beta^s I\left[skill_j = s\right] \times \ln\left(daily\ applications_j\right) + \delta \times I\left[Applications_j > 0\right] + s_j \times t_j + f_j + \epsilon_j \quad (2)$$

In the specification above, we measure labor market tightness as average applications per vacancy day during the first two active weeks for each skill-month. We then interpret the β^s coefficients as the elasticity of posting duration to incremental realizations of application flows controlling for labor market tightness.

applications" group, applications are randomly distributed across vacancy postings (i.e. the conditional distribution of applications is binomial if the overdispersion parameter is zero).

[30] See, for example, Davis, Faberman and Haltiwanger (2012, 2013), Crane et al. (2016), Gavazza, Mongey and Violante (2018), Ketteman, Mueller and Zweimüller (2018), Leduc and Liu (2019) and Mongey and Violante (2019).

Figure 8 plots posting duration elasticity to realized application flows for each skill category. The left axis orders skill categories from lowest to highest mean total application flows. For most skill categories, posting duration declines between 0.24 and 0.41 percent with a one percent increase in daily application flows above the skill category's mean. To place the magnitude of these coefficient in perspective, consider that the mean standard deviation of the natural log of realized daily application flows within skill categories, net of labor market tightness and job function fixed effects, equals 1.1 daily applications. Realized daily application flows across postings, controlling for labor market tightness and job function fixed effects, explain 17.3% of the variation in posting duration.[31]

We next evaluate the contemporaneous and lagged effects of labor market tightness, measured as the average application flows per vacancy, on posting duration. First, we calculate average duration for each skill-time category as the sum of vacancy days (measured in 24 hours intervals) divided by the number of vacancy postings. We measure labor market tightness as the sum of total applications divided by the number of vacancy postings in each skill-time category. Our findings, displayed in Table 9, show that labor market tightness has tiny, negative effects on posting duration. The average effect of realized application flows controlling for labor market tightness, shown in Figure 8, is almost 5 times larger.[32]

A 1 percent increase in average applications per vacancy is associated with a 0.08 percent decrease in posting duration. Heterogeneity across skills drives most of the effect of labor market tightness on application flows. The within-skill R-squared, calculated as the squared correlation of the time-demeaned natural log of mean vacancy duration and the time-demeaned natural log of mean applications per vacancy, is 0.136. Meanwhile, the between R-squared, calculated as the squared correlation of the across-skill mean in the natural log of mean posting duration and the natural log of mean application flows for each month-year is 0.469. This indicates that across-skill heterogeneity, rather than within-skill variation across time, explains most of the variation in vacancy posting duration. As shown in table 9, panel B, the standard deviation in average posting duration within a skill group is small. When comparing vacancy postings across skills, the standard

[31] We calculate the share of the variation in posting duration due to variation in realized daily application flows, net of labor market tightness and job function fixed effects, using the partial R-squared.

[32] The standard deviation of the natural log of mean applications between skillXtime categories is 0.87.

deviation in mean posting duration is almost twice as large. Meanwhile, variation in labor market tightness, measured using the standard deviation in mean application flows, is similar within and across skills.

We interpret these results as indicative of pro-cyclicality in vacancy duration being driven by pro-cyclical screening and recruitment periods, rather than by pro-cyclicality in posting duration. DHI clients' posting duration appears to be insensitive to tightness, measured by mean daily application flows.

5. Application Flows by Job Search Duration

Figure 9 presents the distribution of applicants by total search duration. We measure search duration as total time elapsed from each applicant's first to last application on the DHI platform. While search duration is measured in seconds, we aggregate to 24-hour intervals. To account for different search spells, where an applicant searches for a job for a given period of time, stops (either because they found a job or because they decided to stop searching on the DHI platform), and then begins searching again after a period of inactivity, we also analyze applicants spell (as opposed to overall search) duration. A new spell begins when a job seeker that has had no other applications in the previous 144 hours (60 days) submits a new application. The average (median) applicant has 1.4 (1) search spells. The median spell duration, defined as the time elapsed between each applicant's first and last application within the same search spell, is 35 hours.

Figure 10 shows mean daily applications per active job seeker by search spell duration at the time of application. We define active job seekers as those that submit an application that day or will submit at least one more application within the next 60 days. Job seekers send, on average 3.6 applications during their first 24 hours of search on the platform. Search activity declines over the next days and spikes again each 7 days. The median inactivity period between applications is 9 days. We argue that this pattern is consistent with an optimal response from job seekers to employers' non-sequential search behavior.

6. A Quantitative Sketch of Stages in the Hiring Process

We now draw on our evidence and other studies to create a quantitative sketch of stages in the hiring process. Table 4 says the mean *posting* duration for job openings covered by the DHI Database is 9.4 days. Following Davis, Faberman and Haltiwanger (2013), we use JOLTS data to obtain a mean *vacancy* duration of 40.2 calendar days for job openings in the Information sector. We combine these two pieces of information with evidence from Crane et al. (2016) on the lag

between recruitment events in U.S. labor markets and the start of employment by new hires. Their preferred estimate for the mean value of this start lag is 16.2 days.[33] Figure 7 puts this information together and displays it graphically on a timeline that highlights key events and stages in the hiring process. As indicated in the figure, the total mean elapsed time from date of first posting to the start of employment is 56.2 calendar days.

It's worth stressing that Figure 7 captures only the mean duration of each stage in the hiring process. Our results above uncover great heterogeneity in the duration of posting spells. Likewise, Crane et al. (2016) find great heterogeneity across recruitment events in the length of start lags. Davis, Faberman and Haliwanger (2013) document large differences in mean vacancy durations by industry, employer size, employer growth rate and the employer's worker turnover rate. Thus, Figure 7 is best understood as quantifying average outcomes for phenomena that involve tremendous heterogeneity among employers.

Figure 7. Stages of the Hiring Process and Relationship to Selected Data Sources

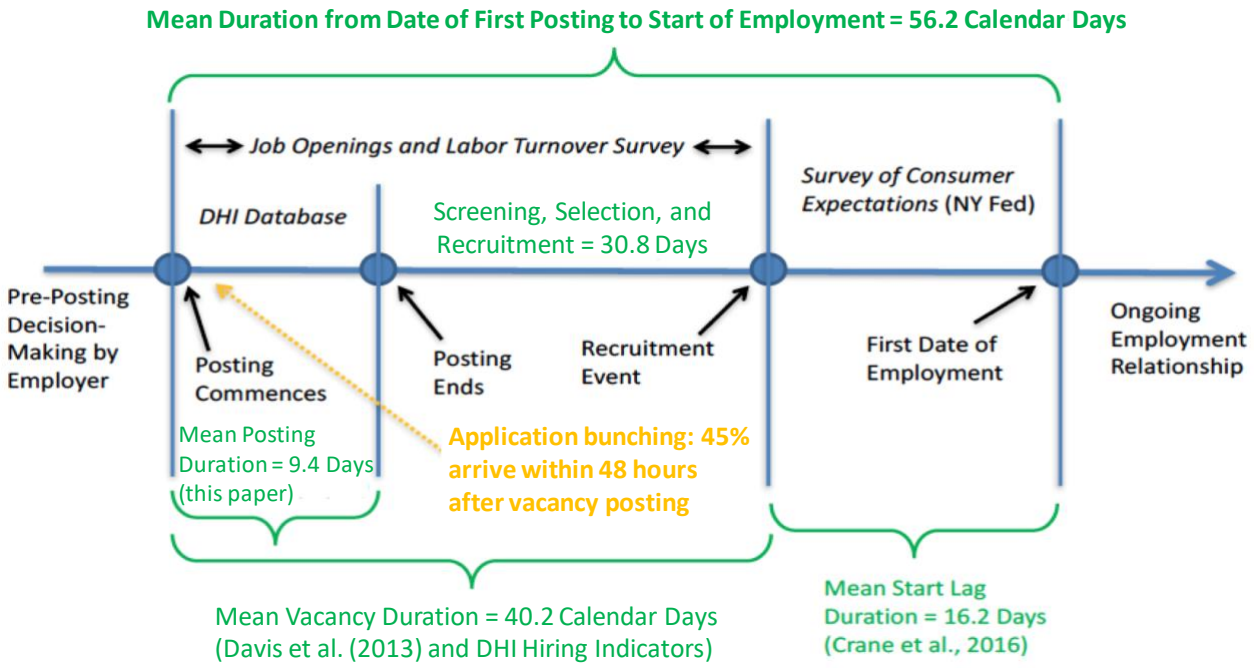

[33] Crane et al. (2016) rely on special supplements to the Federal Reserve Bank of New York's *Survey of Consumer Expectations.* These supplements include recall data from currently employed persons about start lags in their ongoing employment relationships. Crane et al. (2016) do not report evidence specifically for jobs in the Information sector. We make use of their preferred estimate of the mean start lag. Using micro data on German vacancies, Davis et al. (2014) find a mean start lag nearly 40 percent longer than the one obtained for the United States by Crane et al. (2016).

Notes to Figure 7: Mean Posting Duration obtained from the first row in Table 4, which uses data in the DHI Database from January 2012 to December 2017. Mean Vacancy Duration calculated as (7/6) times the average value of the mean vacancy duration for the Information sector from January 2012 to December 2017. Mean Vacancy Duration is calculated using Job Openings and Labor Turnover Survey data and the methodology developed in Davis, Faberman and Haltiwanger (2013). Mean Start Lag is calculated as (7/6) times the preferred estimate of Crane et al. (2016) for the lag between the Recruitment Event and the First Date of Employment. Crane et al. based their estimate on data from the Federal Reserve Bank of New York's Survey of Consumer Expectations. All duration statistics in this figure are expressed in calendars days. The remark in yellow font summarizes a key result in Figures 4 and 5 and Table 5.

No single U.S. data source spans all three stages of the hiring process depicted in Figure 7. That makes it hard to investigate how the various stages in the hiring process relate and interact. In their examination of German micro data, Davis et al. (2014) find evidence that unexpectedly long vacancy durations ("recruitment durations," in their terminology) lead to shorter start lags. Crane et al. (2016) find evidence of countercyclical variation in the mean start lag, a striking contrast to the strongly pro-cyclical behavior of vacancy durations. These results point to important interactions between stages of the hiring process. They also suggest that elapsed time from initial vacancy posting to the start of employment is less pro-cyclical than vacancy durations.

7. Implications for Theories of Search, Matching and Hiring

Figures 4 and 5 and Table 5 above show that job seekers exhibit a striking propensity to direct applications to newly posted vacancies. More than sixty percent of all applications flow to job openings posted within the last 96 hours. This heavy bunching shortly after posting weighs in favor of a non-sequential search strategy, whereby an employer first collects a batch of applications, then screens each applicant in the batch and potentially selects one (or more) for an offer. See Gal, Landsberger and Levykson (1981), Morgan (1983) and Morgan and Manning (1985) on this and other factors that govern the choice between sequential and non-sequential search strategies.

One might interpret our evidence that job seekers strongly favor newly posted vacancies as supporting the empirical relevance of stock-flow matching models, as in Coles and Smith (1998) and Gregg and Petrongolo (2005). However, we also find that the mean posting duration in the DHI Database is only about one-fifth as long as the mean vacancy duration for comparable jobs in the JOLTS. A quarter of the postings have completed spell durations of 2.9 days or less (Table 4). We see this combination of results as hard to reconcile with sequential search by employers – or, at least for a large share of the employers covered by our data.

We are not the first to argue that hiring behavior is, for many employers, inconsistent with sequential search. In data for Dutch employers, Van Ours and Ridder (1992) find that almost all hires take place from a pool of applicants formed shortly after vacancy posting. They also find that the vacancy filling rate is low for new vacancies and rises with vacancy age. These findings are broadly consistent with our evidence and the Figure 7 sketch of stages in the hiring process. That sketch fits well with non-sequential employer search behavior, not so well with sequential search. Van Ommeren and Russo (2008) formally reject the hypothesis of sequential search by Dutch employers who rely on advertising or employment agencies to recruit workers, which constitute nearly half the hires in their sample. They do not reject sequential search for other types of hires covered by their study.

As Morgan and Manning (1985) and Gautier (2002) point out, because non-sequential employer search creates a delay between the submission of an application and the employer's selection of a recruit, it creates incentives for job seekers to adopt a non-sequential search strategy as well. It makes sense for job seekers to submit applications for multiple job openings while awaiting call-backs and offers, unless applications themselves are very costly to submit. Abbring and Van Ours (1994) provide evidence that job seekers behave in this manner. Our results from the applicants' perspective (in section III.7) are also consistent with non-sequential search strategies.

Despite systematic and casual evidence in its favor, the non-sequential perspective has been thoroughly overshadowed by theories in the mold of Diamond (1982), Mortensen (1982), Pissarides (1985) and Mortensen and Pissarides (1994). These theories postulate sequential search by employers and workers. In the past twenty years or more, the leading treatments of frictional unemployment, job-finding rates, vacancy behavior, labor market tightness, wage dispersion with search frictions, and job creation incentives in environments with search frictions have been entirely dominated by the sequential search perspective. In addition to the seminal contributions just cited, leading examples include Burdett and Mortensen (1998), Pissarides (2000), Postel-Vinay and Robin (2002), Mortensen (2003), Hall (2005), Shimer (2005), Hornstein, Krusell and Violante (2011) and Davis, Faberman and Haltiwanger (2013). It strikes us as potentially problematic to rely on sequential search models for detailed quantitative assessments of labor market outcomes and policy interventions, when the characterization of the hiring process in these models is so sharply at odds with actual hiring behavior.

Theories of non-sequential search date to Stigler (1961). Gal, Landsberger and Levykson (1981), Morgan (1983) and Morgan and Manning (1985) theoretically analyze the choice between sequential and non-sequential search strategies. Labor market environments with non-sequential search involve a different set of externalities than environments characterized by sequential search. See Gautier (2002), Albrecht, Gautier and Vroman (2006) and Wolthoff (2014) on this point. Wolthoff (2014) consider a theoretical model with two-sided search with employers and job seekers who can simultaneously contact multiple parties on the other side of the labor market. We see our evidence as strong motivation for attention to non-sequential search by employers, as well as models in which both job seekers and employers can simultaneously contact multiple potential partners with whom to initiate an employment relationship.

IV. The Effects of Changes in Platform-Level Search and Matching Functionality

On December 14, 2014, Dice.com launched a major update to its platform. On the job seeker side, the changes made it easier for users to register and create a profile. The updated platform also included a simplified application process and an improved search engine to help job seekers efficiently customize their search.[34] From the employer side, the update offered DHI clients the option to search through candidates' profiles and directly contact potential hires to encourage them to apply. The changes in Dice.com's functionality are in line with DHI's broader goal of providing "customized search engines and audience-tailored websites (that) are efficient and relevant, easy to use and valuable to our users, helping us build a loyal and engaged audience." (Dice Holdings Inc. 10-K filing for the fiscal year ended December 31, 2014).

Figure 11 provides evidence consistent with the platform change leading to an increase in the average number of applications per applicant, and a decrease in the time elapsed between applications. The change is particularly pronounced for Email applications. The number of total applications and the average number of applications per applicant rose sharply after December 2014 (Panels A and B). Consistent with the platform change streamlining the process to submit Email applies, the rise is driven by application flows to job postings that allow this channel. To further analyze whether applying became easier after the platform change, for each job seeker, we calculate the time elapsed between each application they submit, separately for email and URL applications. We define "high-frequency" applications as those an applicant submits within a 60-second interval

[34] We provide additional information on the platform and its evolution after 2014 in Appendix B.

of their previous application through the same channel (email or URL). Analyzing the evolution of high-frequency application allows us to examine applicants' propensity to bunch applications together in very short time intervals. Panel D indicates that after the platform update on December 2014 applicants submitted a larger share of email applications in quick succession.

In Figure 12, we examine the effects of the platform change on the distribution of applications by employer size.[35] The effect of the platform change is large and decreasing in firm size. Applications per vacancy day rose sharply, particularly for the 3 smallest employer size categories: 1 to 9 employees, 10 to 19, and 20 to 99 employees. These results highlight the potential of online job boards to affect not just the volume, but also the allocation of applicants across employers.

V. Concluding Remarks

To Be Written

[35] We focus on Direct Hire job postings since it is less clear what "employer size" refers to for Recruitment or Staffing Firms.

References

Abbring, J.H., and J.C. Van Ours, 1994. "Sequential or Non-Sequential Employers' Search?" *Economics Letters,* 44, 323-328.

Albrecht, James, Bruno Decreuse and Susan Vroman, 2019. "Directed Search with Phantom Vacancies," working paper, Georgetown University.

Albrecht, James, Pieter A. Gautier and Susan Vroman, 2006. "Equilibrium Directed Search with Multiple Applications," *Review of Economic Studies,* 73, 869-891.

Anderews, Martyn, Steve Bradley, Dave Scott and Richard Upward, 2017. "Estimating the Stock-Flow Matching Model Using Micro Data," *Journal of the European Economic Association,* 11, no. 5 (October), 1153-1177.

Arnosti, ick, Ramesh Johari and Yash Kanoria, 2015. "Managing Congestion in Matching Markets," available at semanticscholar.org.

Autor, David H., 2009. "Introduction to Studies of Labor Market Intermediation*,"* in *Labor Market Intermediaries,* edited by David H. Autor, University of Chicago Press.

Bailey, Martin J., Richard F. Muth and Hugh O. Nourse, 1963. "A Regression Method for Real Estate Price Index Construction," *Journal of the American Statistical Association,* 58, 933-942.

Banfi, Stefano and Benjamin Villena-Roldan, 2018. "Do High-Wage Jobs Attract More Applicants? Directed Search Evidence from the Online Labor Market," forthcoming, *Journal of Labor Economics.*

Baydur, Ismail, 2017. "Worker Selection, Hiring and Vacancies," *American Economic Journal: Macroeconomics,* 9, no. 1, 88-127.

Barron, John M., Mark C. Berger and Dan A. Black. 1997. "Employer Search, Training, and Vacancy Duration," *Economic Inquiry,* 35 (January), 167-192.

Barron, John M., John Bishop and William C. Dunkelberg, 1985. "Employer Search: The Interviewing and Hiring of New Employees," *Quarterly Journal of Economics,* 67, 43-52. "...most employment is the outcome of an employer selecting from a pool of job applicants..."

Belot, Michele, Philipp Kircher and Paul Muller, 2018. "How Wage Announcements Affect Job Search – A Field Experiment," Working Paper in Economics No. 739, University of Gothenberg.

Brencic, Vera and John B. Norris, 2012. "Employers' On-Line Recruitment and Screening Practices," *Economic Inquiry,* 50, no. 1 (January), 94-111.

Bull, Clive, Oscar Ornati and Piero Tedeschi, 1987. "Search, Hiring Strategies, and Labor Market Intermediaries," Journal of Labor Economics, 5, no. 4, part 2 (October), S1-S17.

Burdett, Kenneth and Dale T. Mortensen, 1998. "Wage Differentials, Employer Size, and Unemployment," *International Economic Review,* 39, no. 2 (May), 257-273.

Buss, Jason and Kasia Senderowska, 2015. "Right Place, Right Time: The Data behind Hiring Success," at https://www.smartrecruiters.com/blog/right-place-right-time-the-data-behind-hiring-success/.

Butters, Gerard R., 1977. "Equilibrium Distributions of Sales and Advertising Prices," *Review of Economic Studies,* 44, 465-491.

Carrillo-Tudela, Carlos and William Hawkins, 2017. "Stock-Flow Marching with Heterogeneous Workers and Firms: Theory and Evidence from the U.K." Slide deck, 7 December 2017.

Carrillo-Tudela, Carlos, Hermann Gartner and Leo Kaas, 2018. "Understanding Vacancy Yields: Evidence from German Data," work in progress, June.

Case, Bradford and John M. Quigley, 1991. "The Dynamics of Real Estate Prices," *Review of Economics & Statistics,* 73, no. 1 (February), 50-58.

Case, Karl E. and Robert J. Shiller, 1987. "Prices of Single Family Homes Since 1970: New Indexes for Four Cities," *New England Economic Review,* September-October, 45-56.

Case, Karl E. and Robert J. Shiller, 1989. "The Efficiency of the Market for Single-Family Homes," *American Economic Review,* 79, no. 1 (March), 125-137.

Cheron, Arnaud and Bruno Decreuse, 2017. "Matching with Phantoms," *Review of Economic Studies,* 84, no. 3 (July), 1041-1070.

Chugh, Sanjay K. and Christian Merkl, 2016. "Efficiency and Labor Market Dynamics in a Model of Labor Selection," *International Economic Review,* forthcoming,

Coles, Melvyn and Eric Smith. 1998. "Marketplaces and Matching," *International Economic Review,* 39, 239-255.

Court, Andrew, 1939. "Hedonic Price Indexes with Automotive Examples," in *The Dynamics of Automobile Demand,* General Motors, New York.

Crane, Leland, Jr., Steven J. Davis, R. Jason Faberman, and John C. Haltiwanger, 2016. "A Generalized Model of Hiring Dynamics: Evidence, Implementation, Implications," working paper in progress, University of Chicago.

Davis, Steven J., 1990. "Size Distribution Statistics from County Business Patterns Data," unpublished paper, University of Chicago, September.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger, 2012. "Recruiting Intensity during and after the Great Recession: National and Industry Evidence," *American Economic Review Papers & Proceedings,* 102, no. 3, 584–588.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger, 2013. "The Establishment-Level Behavior of Vacancies and Hiring," *Quarterly Journal of Economics,* 128, no. 2, 581–622.

Davis, Steven J., John Haltiwanger, Ron Jarmin and Javier Miranda, 2007. "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms," *NBER Macroeconomics Annual,* 21, 107-180.

Davis, Steven J., Christof Röttger, Anja Warning and Enzo Weber, 2014. "Recruitment and Vacancy Durations in Germany," University of Regensburg Working Papers in Business, Economics and Management Information Systems, Working Paper No. 481, 14 June.

Davis, Steven J. and Brenda Samaniego de la Parra, 2016. "The DHI Vacancy and Application Flow Database: Record Layouts, Variable Descriptions, and Basic Summary Statistics," DHI Database Technical Memorandum Number 1.

DHI Group, Inc., 2016. *Form 10-K, 2016* at https://dhigroupinc.com/investors/default.aspx.

Deming, David and Lisa B. Kahn, 2018. "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals," *Journal of Labor Economics,* 36, no. S1 (January).

DHI Hiring Indicators, 2016. "Mean Job Vacancy Duration Fell to 28.3 Working Days in June," Report 28, August, at http://dhihiringindicators.com.

Diamond, Peter A., 1982. "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy,* 90, no. 5, 881-894.

Eeckhout, Jan and Phlippp Kircher, 2010. "Sorting versus Screening: Search Frictions and Competing Mechanisms," *Journal of Economic Theory,* 145, no. 4 (July), 1354-1385.

Faberman, R. Jason and Marianna Kudlyak, 2016. "The Intensity of Job Search and Search Duration," Working Paper 14-12R (February), Federal Reserve Bank of Richmond.

Gal, S., M. Landsberger and B. Levykson, 1981. "A Compound Strategy for Search in the Labor Market," *International Economic Review,* 22, no. 3 (October), 597-608.

Galenianos, Manolis and Phlipp Kircher, 2009. "Directed Search with Multiple Job Applications," *Journal of Economic Theory,* 144, 445-471.

Gavazza, Allesandro, Simon Mongey and Gianluca Violante, 2018. "Aggregate Recruiting Intensity," *American Economic Review,* 108, no. 8 (August).

Gautier, Pieter A., 2002. "Non-sequential Search, Screening Externalities and the Public Good role of Recruitment Offices," *Economic Modelling,* 19, 179-196.

Gautier, Pieter A. and Christian L. Holzner, 2016. "Simultaneous Search and Efficiency of Entry and Search Intensity," working paper.

Gregg, Paul and Barbara Petrongolo, 2005. "Stock-Flow Matching and the Performance of the Labor Market," *European Economic Review,* 49, no 8 (November), 1987-2011.

Goodman, Allen C., 1998. "Andrew Court and the Invention of Hedonic Price Analysis," *Journal of Urban Economics,* 44, no. 2 (September), 291-298.

Griliches, Zvi, 1971. "Hedonic Price Indexes Revisited," in *Price Indexes and Quality Change,* edited by Z. Griliches, Harvard University Press: Cambridge, MA.

Hall, Robert E., 2005, "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review,* 95(1), 50-65.

Hornstein, Andreas, Per Krusell and Giovanni L. Violante, 2011. "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," *American Economic Review,* 101, no. 7 (December), 2873-2898.

Kettemann, Andreas, Andreas I. Mueller and Josef Zweimüller, "Vacancy Durations and Entry Wages: Evidence from Linked Vacancy-Employer-Employee Data," NBER Working Paper No. 25118.

Kolchin, Valentin F., 1978. *Random Allocations.* Washington, Winston.

Kotz, Samuel and Norman Lloyd Johnson, 1977. *Urn Models and Their Applications.* New York: John Wiley & Sons.

Kudlyak, Marianna and Jason Faberman, 2016. "The Intensity of Job Search and Search Duration," Federal Reserve Bank of San Francisco, Working Paper 2016-13.

Kudlyak, Marianna, Damba lkhagvasuren and Roman Sysuyev, 2014. "Systematic Job Search: New Evidence from Individual Job Application Data," working paper, September. See message from Marianna that points out connections to work.

Kruger, Alan B. and Andreas Mueller, 2010. "Job Search and Unemployment Insurance: Evidence from Time Use Data," *Journal of Public Economics,* 94, 298-307. They state (page 300) that job search is heavily concentrated on weekdays and report "Nearly a quarter of the unemployed engage in job search activities on any given weekday, compared with 6.7% on weekends."

Leduc, Sylvain and Zheng Liu, 2019. "The Weak Job Recovery in a Macro Model of Search and Recruiting Intensity," *American Economic Journal: Macroeconomics,* forthcoming.

Manning, Alan, 2000. "Pretty Vacant: Recruitment in Low-Wage Labour Markets," *Oxford Bulletin of Economics and Statistics,* 62, Special Issue, 747-770.

Marinescu, Iona and Roland Rathelot, 2015. "Mismatch Unemployment and the Geography of Job Search," working paper, University of Chicago.

Marinescu, Iona and Ronald Wolthoff, 2015. "Opening the Black Box of the Matching Function: The Power of Words," IZA Discussion Paper No. 9071.

Mongey, Simon and Gianluca Violante, 2019. "Macro Recruiting Intensity from Micro Data," NBER Working Paper 26231.

Morgan, Peter B., 1983. "Search and Optimal Sample Sizes," *Review of Economic Studies,* 50, no. 4, 659-675.

Morgan, Peter B. and Richard Manning, 1985. "Optimal Search," *Econometrica,* 53, no. 4 (July), 923-944.

Mortensen, Dale T., 1982. "The Matching Process as a Noncooperative Bargaining Game," in John J. McCall, editor, *The Economics of Information and Uncertainty,* University of Chicago Press.

Mortensen, Dale T., 2003. *Wage Dispersion: Why Are Similar Workers Paid Differently?* Cambridge, MA.

Mortensen, Dale T. and Christopher A. Pissarides, 1994. "Job Creation and Job Destruction and the Theory of Unemployment," *Review of Economic Studies,* 61, 397-415.

Moylan, S., J. Millar and R. Davis, 1982. "Unemployment – The Year After," Department of Employment, Employment Gazette, 90, 334-340.

Pissarides, Christopher A., 1979. "Job Matching with State Employment Agencies and Random Search," *Economic Journal,* 89, 818-833.

Pissarides, Christopher A., 1985. "Short-run Equilibrium Dynamics of Unemployment, Vacancies and Real Wages," *American Economic Review,* 75, no. 4, 676-690.

Pissarides, Christopher A., 2000. *Equilibrium Unemployment Theory,* second edition. Cambridge, MA: MIT Press.

Postel-Vinay, Fabien and Jean-Marc Robin, 2002. "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica,* 70, no. 6 (November), 2295-2350.

Russo, Giovanni, Wolter H.J. Hassink and Cees Gorter, 2005. "Filling Vacancies: An Empirical Analysis of the Cost and Benefit of Search in the Labour Market," *Applied Economics,* 37, 1597-1606.

Russo, Giovanni, Piet Rietveld, Peter Nijkamp and Cees Gorter, 2000. "Recruitment Channel Use and Applicant Arrival: An Empirical Analysis," *Empirical Economics,* 25, 673-697.

Shimer, Robert, 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review,* 95, no. 1, 25-49.

Skandalis, Daphne, 2018. "Breaking News: Information about Firms' Hiring Needs Affects the Direction of Job Search," working paper, https://economics.ceu.edu/sites/economics.ceu.edu/files/attachment/event/1117/jmpdaphneskandalis.pdf

Stigler, George J., 1961. "The Economics of Information," *Journal of Political Economy,* 69, no. 3 (June), 213-225.

Van Ommeren, Jos and Giovanni Russo, 2008. "Firm Recruitment Behavior: Sequential or Non-sequential Search?" *Oxford Bulletin of Economic Statistics,* 76, no. 3, 432-455.

Van Ours, Jan and Geert Ridder, 1992. "Vacancies and the Recruitment of New Employees," *Journal of Labor Economics,* 4, no. 1, 138-155.

Van Ours, Jan and Geert Ridder, 1993. "Vacancy Durations: Search or Selection?" *Oxford Bulletin of Economics and Statistics,* 55, no. 2 (May), 187-198.

van de Ven, R., and N.C. Weber, 1993, "Bounds for the median of the negative binomial distribution" Metrika 40, 185–189

Weber, A., 2000. "Vacancy Durations: A Model of Employer Search," *Applied Economics,* 32, 1069-1075.

Wolthoff, Ronald, 2018. "Applications and Interviews: Firms' Recruiting Decisions in a Frictional Labor Market," *Review of Economic Studies*, 85, no. 2, 1314-1351.

Table 1. Vacancy Postings and Applications in the DHI Database, January 2012 to December 2017

	Total Millions	Direct Hire Millions	Recruitment and Staffing Firm Millions
(1) Number of Raw Vacancy Postings	7.5	2.5	5.0
(1.a) Standard Vacancy Postings	5.6	1.7	3.9
(1.b) Long-Duration Vacancy Postings	1.9	0.8	1.1
(2) Number of Vacancies, After Slicing the Long-Duration Postings	11.7	4.3	7.4
(3) Volume of Applications	125.3	47.9	77.4
(3.a) Email Applications	95.3	34.4	60.9
(3.b) URL Applications	30.0	13.4	16.6
(4) Volume of Non-Duplicate Applications	114.4	43.4	71.0
(4.a) Email Applications	88.9	32.1	56.8
(4.b) URL Applications	25.5	11.3	14.2

Notes: "Direct Hire" clients seek to hire their own employees, "Recruitment" firms solicit applicants for third parties, and "Staffing" firms hire workers to lease to other firms. Row (1) pertains to the number and distribution of distinct Job ID values in the DHI Database ("*Activity File")* and Row (2) pertains to Vacancy ID values (*"Activity File"*). See text for an explanation of "Long-Duration Postings" and how we slice them to construct Vacancy IDs. Row (3) reports the number and distribution of applications (from the *"Detailed Applications File"*). "Email Applications" refer to those submitted through DHI, and "URL Applications" refer to the frequency with which job seekers click through to an external URL. Row (4) presents the count of non-duplicate applications. Duplicate applications are those submitted by the sample applicant ID to the same job ID more than once.

Table 2. Distribution of Direct Hire Vacancies with Positive Applications by Employer Size

0 Employees	1-4	5-9	10-19	20-49	50-99	100-249
18.2%	12.7%	5.7%	5.7%	11.0%	8.1%	7.5%
250-499	500-999	1,000-2,499	2,500-4,999	5,000-9,999	10,000+	
6.2%	2.8%	4.1%	3.2%	3.1%	11.6%	

Notes: In constructing this table, each Vacancy ID with one or more applications receives equal weight, and Vacancy IDs with no applications receive zero weight. The distribution of vacancies by employer size pertains to privately held and publicly listed companies. Employer size is obtained from Dunn & Bradstreet, typically when the client opens a new account and may not be current.

Table 3. Summary Statistics for Frequently Posted Job Titles in the DHI Database

(1) Minimum Posting Frequency	(2) Number of Job Titles	(3) Share of Job IDs	(4) Share of Vacancy IDs	(5) Share of Applications
250 Job IDs	1,285	93.5%	94.0%	95.2%
100 Job IDs	1,983	95.0%	95.5%	96.5%
50 Job IDs	2,746	95.7%	96.2%	97.1%

Notes: Column (2) reports the number of distinct job titles that meet the minimum posting frequency specified in Colum (1). Columns (3) to (5) report the shares of Job IDs, Vacancy IDs and Applications accounted for by these frequently posted job titles.

Table 4. The Distribution of Completed Posting Durations by Job Type and Applications Volume

			Percentile				
	No. of Postings	Mean	10	25	50	75	90
All Standard Postings	5,362,717	9.4	1.0	2.9	7.0	14.0	22.7
All Job Titles with at Least 100 Standard Postings	5,139,696	9.4	1.0	2.9	7.0	14.0	22.6
Selected Job Types							
Developer	1,181,708	8.9	1.0	2.3	6.8	13.9	21.5
Engineer	626,241	10.7	1.1	3.8	7.4	16.0	25.0
Administrator	388,857	9.0	1.0	2.6	6.8	13.9	21.9
Mechanical Engineer	6,133	11.5	1.4	4.3	9.0	17.0	26.2
Electrical Engineer	6,010	12.2	2.0	5.0	10.1	18.5	27.0
Business Analyst	226,768	8.9	1.0	2.7	6.9	13.2	21.8
Analyst	326,291	10.0	1.0	3.1	7.0	14.9	23.9
Help / Support Desk	246,829	10.0	1.1	3.2	7.0	15.0	22.9
Sales / Business Development	35,043	11.2	1.0	3.5	8.5	17.0	26.0
By Number of Applications							
No Application	1,092,895	6.1	1.0	1.3	4.0	7.4	15.0
1 Application	740,529	7.4	1.0	2.0	5.7	10.0	18.0
2-4 Applications	1,269,761	9.3	1.0	3.1	7.0	13.9	21.2
5-9 Applications	910,746	11.1	1.7	4.7	8.1	16.8	25.0
10-19 Applications	656,076	12.2	1.8	5.0	10.0	19.0	26.9
20+ Applications	692,710	12.2	1.1	4.7	10.0	19.8	27.0
N.B. Using Elapsed Time Net of Offline Spells, All Standard Postings	5,362,717	9.1	1.0	2.8	6.9	13.9	21.8

Notes: Table entries report statistics on completed spell durations for standard vacancy postings from January 2012 to November 2017. We measure duration from initial posting date-time to final removal date-time in seconds and express the statistics in 24-hour intervals. The bottom row considers an alternative duration measure that nets out offline spells. For example, if a vacancy is first posted for 48 hours, taken offline for 24 hours, and then reposted for 72 hours prior to permanent removal, the alternative vacancy duration measure is 48 + 72 hours, which amounts to 5.0 days. In constructing this table, we dropped observations with first posting date on or after December 1, 2017 to avoid the inclusion of incomplete spells.

Table 5. Applications Per Vacancy and Application Bunching at Young Postings

	Mean Number of Applications Per Vacancy		*Applications in First 30 Days Since Posting, Percent Received Within*			
			First 48 Hours After Posting		First 96 Hours After Posting	
	(1)	(2)	(3)	(4)	(5)	(6)
	EW	FW	EW	FW	EW	FW
All Standard Vacancy Postings	11.0	88.1	46.7	45.3	62.9	59.9
All Job Titles with at Least 100 Standard Vacancy Postings	11.2	89.1	46.8	45.3	62.9	60.0
Sorted by Applications Volume						
Bottom Quintile	3.4	19.8	37.4	32.9	54.5	48.3
Fourth Quintile	5.4	30.7	42.3	38.1	58.6	52.9
Third Quintile	7.4	41.9	44.0	38.6	60.5	53.7
Second Quintile	10.8	60.4	49.2	44.8	65.2	59.7
Top Quintile	22.3	134.6	52.1	49.8	67.5	64.0
Selected Job Categories						
Developer	16.3	141.3	50.1	49.9	65.6	64.2
Engineer	7.5	64.4	39.6	40.9	56.1	55.7
Administrator	10.4	58.5	48.7	45.3	64.7	60.2
Mechanical Engineer	4.1	17.5	35.3	26.0	51.2	41.1
Electrical Engineer	3.7	15.2	31.3	24.4	49.1	40.6
Business Analyst	22.5	97.0	50.2	49.5	66.3	63.1
Analyst	9.9	67.4	43.0	39.5	59.6	54.4
Help / Support Desk	7.5	32.5	39.2	29.7	56.9	45.4
Sales / Business Development	3.0	24.0	32.8	28.5	49.6	43.6

Notes: Except for the first row, all table entries pertain to standard vacancy postings with at least 100 distinct postings. We assign equal weight to each vacancy in columns headed "EW" and weight by the flow of applications received in columns headed "FW". To sort by applications volume, we first compute the mean applications per vacancy at the job title level, and we then sort job titles into quintiles based on mean applications per vacancy. In constructing this table, we dropped observations for the last month in the sample to avoid the inclusion of incomplete spells. Columns (3) to (6) include only postings that receive at least one application.

Table 6. The Distribution of Application Counts over Vacancy Postings, Summary Statistics and Comparisons to Random Assignment and Zero Inflated Negative Binomial Models with dispersion parameters that match the data's standard deviation and share of zero applications, and simple mean and standard deviation respectively.

	Direct Hire Standard Postings			
	(1)	(2)	(3)	(4)
	DHI Data	Model mean (μ), overdispersion (θ), inflated zero probability (p)		
		Binomial (i.e. Random Assignment)	Negative Binomial	Zero-Inflated Negative Binomial
		(μ=11.8)	(μ=11.8, θ=5.8)	(μ=18.9, θ=2.2, *p*=0.005)
Simple Mean	11.8	11.8	11.8	18.8
Standard Deviation	28.6	3.4	28.6	28.6
Percent with 0 applications	18.9	0.0008	48.1	18.9
Percent with 1 application	12.1	0.0089	8.2	8.1
Flow-Weighted Mean	81.1	12.8	81.1	62.2
Ratio of Flow-Weighted to Simple Mean	6.9	1.1	6.9	3.3
Ratio of Median to Simple Mean	0.34	1	0.08	0.39

Notes: The columns headed "DHI Data" report values for the data depicted in Figure 4 (Direct Hire and Recruitment Firm standard postings). The columns headed "Random Assignment" report model-implied values when applications are distributed randomly and independently to postings. In the model, a given application flows to a particular posting with a probability of $1/v$.

The columns headed ZINB(α) report the values implied by zero negative binomial model with dispersion parameter equal to α, and conditional mean for the negative binomial equal to the observed sample mean divided by 1 minus the probability of being in the zero-application case. In a ZINB model, the share of vacancy postings with zero applications is equal to $p+(1-p)(1+\mu\alpha)^{-1/\alpha}$ where p is the probability of the job posting being in the "always zero applications" group (i.e. the "inflated zero" group), μ is the mean for the negative binomial component of the model (i.e. the conditional mean for applications given the job posting does not belong to the "always zero applications" group) and α is the dispersion parameter.

For this table, we selected *p* and μ to target the observed sample mean and the share of job postings with zero applications. For Direct Hires, the probability of being in the "always zero applications" group is 13.0% and 1.2% for the model with overdispersion parameter equal to 1 and 2, respectively. The analogous probabilities for Recruitment Firms are 15.9% and 1.2%.

To calculate the median for the zero-inflated negative binomial model, we simulated 100,000 observations with each distribution. Van de Ven and Weber (1993) obtain bounds for the median for a negative binomial model as a function of the distribution's mean and overdispersion parameters. The values we obtained for the simulated median of a zero-inflated negative binomial distribution lie within these bounds.

Table 7. How Application Counts Vary by Wage Posting, Permissibility of Third-Party Applications, Client Type, Applications Channel, and Unobserved Heterogeneity

Unit of Analysis: Completed Standard Vacancy Posting Spells
Dependent Variable: total application count
Model: Zero-inflated negative binomial

Controls	Wage Posting Dummy	3rd party OK Dummy	Client Type Direct Hire ("RF" baseline)	Email Application Dummy	Dispersion Parameter
Baseline, with 55 Skill FE	0.03 (0.002)	1.02 (0.001)	0.26 (0.001)	0.19 (0.002)	2.04 (0.001)
+ 72 Time FE	0.07 (0.002)	0.87 (0.001)	0.29 (0.001)	0.10 (0.001)	1.88 (0.001)
+ 15 Client Size Categories	0.09 (0.002)	0.85 (0.001)	0.28 (0.001)	0.05 (0.002)	1.86 (0.001)
+ 54 Job Location (50 States, PR, VI, DC and Other)	0.08 (0.002)	0.87 (0.001)	0.28 (0.001)	0.03 (0.002)	1.81 (0.001)
+ 1,713 Job Title FE	0.06 (0.002)	0.87 (0.001)	0.33 (0.001)	0.05 (0.002)	1.61 (0.001)
+ 4,408 Employer FE	0.02 (0.002)	1.02 (0.002)	0.36 (0.005)	0.22 (0.003)	1.33 (0.001)

Notes: The sample includes standard postings where the job description mentions the skill category listed in the first column as the first required skill. The sample period is January 2012 to November 2017. We exclude postings with first posting date on or after December 1, 2017 to limit the sample to completed spells. All specifications include a set of 55 skill requirements fixed effects.

Moving downward from the row headed "Time Fixed Effects", each successive row adds the indicated fixed effects to the controls for the negative binomial component of the ZINB model.

In the zero-inflated component of the model, we use an indicator variable equal to 1 for postings with zero views as the control to identify the "always zero applications" job postings.

For each vacancy posting, DHI clients can choose whether to accept applications from 3rd party applicants, and this decision can change throughout the duration of a posting. For this table, we classify postings as "3rd party OK" if they accept 3rd party applications for the majority of their posting duration. The majority of standard vacancy postings keep the same decision regarding allowing 3rd party applicants throughout their duration. In the DHI data, 4 percent of all vacancy-days have a value of "Unknown" for the variable that identifies whether the posting accepted third party applications or not. For this table, we assume that vacancy postings that do not explicitly authorize third party applications do not allow them. Hence, the "3rd party OK" indicator variable in this table is equal to zero if the data reports a value of "Unknown."

Vacancy postings can accept applications through only one channel (Email or URL) at any given point in time. (See *The DHI Vacancy and Application Flow Database: Record Layouts, Variable Descriptions, and Summary Statistics* for further information on the DHI data.) However, DHI clients can choose to change the channel of application for each of their vacancy postings through time. We can only determine the selected application channel by looking at the channel through which applications arrive for each posting. We classify a posting as accepting "Email applications" if most of its applications are via Email and as URL if most of the applications are through this channel. The majority of postings (with applications) accept only one application channel throughout their

duration. We cannot observe the selected application channel for vacancy postings that receive zero applications. For this table, we randomly assign zero application postings to an application channel. Most vacancy postings accept only one type of applications throughout their entire posting duration. We report standard errors in parenthesis.

Table 8. Changes in Overdispersion Estimates as we account for additional job posting heterogeneity

Unit of Analysis: Completed Standard Vacancy Posting Spells
Dependent Variable: total application count
Model: Zero-inflated negative binomial

			Applications per Vacancy		*Overdispersion Estimates (Maximum likelihood)*									
Skill Category	# of Postings	Share of Postings with No Applications (%)	Mean	Std. Dev.	Intercept Only Model	+Posting Duration (Hours)	+ Time (72)	+ Direct Hire Dummy	+ Employer Size Category + Real Wage (15) + (3)	+ 3rd party OK? (3)	+ State FE (54)	+Email Channel	+ Job titles FE	+Employer FE
APPLICATION	103,777	26%	5.4	11.4	2.04	1.76	1.65	1.63	1.56	1.44	1.35	1.35	1.17	0.94
DATA	135,185	15%	10.9	21.0	1.90	1.78	1.44	1.44	1.38	1.13	1.09	1.08	0.90	0.74
JAVA	306,931	16%	21.3	60.6	2.90	2.85	1.87	1.85	1.80	1.27	1.24	1.24	1.12	0.94
.NET	17,711	15%	13.2	28.3	2.22	2.17	1.38	1.35	1.26	0.88	0.83	0.83	0.78	0.49
ORACLE	183,322	14%	12.3	24.0	1.86	1.73	1.30	1.30	1.29	1.16	1.13	1.13	0.92	0.76
SAP	195,610	11%	11.5	17.4	1.42	1.31	1.12	1.12	1.11	1.01	1.00	1.00	0.95	0.79
NETWORK	161,689	22%	9.6	25.3	2.60	2.52	1.93	1.88	1.70	1.32	1.28	1.28	1.18	0.95
SECURITY	97,219	29%	4.6	9.7	2.03	1.73	1.50	1.48	1.37	1.17	1.13	1.12	0.95	0.73
SOFTWARE	194,161	26%	6.5	16.0	2.32	2.12	1.94	1.94	1.85	1.65	1.58	1.58	1.34	1.07
SQL	103,261	11%	18.5	38.5	2.02	1.93	1.46	1.44	1.35	0.95	0.92	0.92	0.86	0.69
SYSTEMS	250,817	23%	7.3	15.7	2.13	1.98	1.74	1.73	1.62	1.37	1.32	1.32	1.09	0.89
WEB	98,021	26%	8.1	24.2	2.87	2.80	2.01	1.98	1.86	1.52	1.47	1.47	1.31	1.04

Notes to Table 8: The sample includes standard postings where the job description mentions the skill category listed in the first column as the first required skill. The sample period is January 2012 to November 2017. We exclude postings with first posting date on or after December 1, 2017 to limit the sample to completed spells.

The "Intercept-only model" controls for positive views in the zero-inflated part of the model and only allows for an intercept in the negative binomial part of the model.

Moving rightward from the column headed "Intercept-only model," each successive column adds the indicated variable to the regressor list in the negative binomial part of the model. We report the unit of measure or the number of distinct categories in parentheses.

We observe posting duration, first active date (used for time fixed effects), employer and employer type for all postings. However, not all standard postings have information for city of employment, employer size, job title, and pay rate. To maintain the same sample size across the various columns, we include a "missing" category for those regressors with incomplete information.

For "State Effects" we include 54 fixed effects for all 50 states, plus Puerto Rico, Virgin Islands, and DC. We also include an "Other" category for locations outside the US.

For "Employer Size + Real Wage" we add the 15 size categories listed in Table 3 Panel B, interacted with a dummy for Direct Hire employers. We then also include 3 wage indicators for postings that offer an explicit, numerical hourly, monthly, or annual wage offers.

We interact each of these dummies with the natural log of the real wage. When the posting specifies a compensation range, we use the midpoint. We measure wages in 2016 dollars using the GDP implicit price deflator. 80 percent of standard postings include some information regarding pay rate. However, only 22 percent of these postings with non-missing pay rate information (18 percent of all standard postings) include a numerical wage (the rest make a statement about pay being based on experience, competitive, market based, etc.). We group the employer size effect with the wage effect because we find that, after controlling for employer size, wage does not change the overdispersion estimates.

See the notes in table 7 regarding the "3rd party OK" and "Email Channel" indicator variables.

Each vacancy posting is assigned to exactly one extended job title. An extended job title is the extended job title which is composed by seniority level + skill requirement + job function. When adding controls for "job titles," we include a separate fixed effect for each extended job title with at least 100 distinct postings. All other job titles are grouped into an "other" category for this analysis. The number of job title fixed effects varies from 6 to 89, depending on the skill category.

Table 9. Posting Duration as a Function of Tightness (Measured Using Mean Application Flows per Vacancy), Monthly Data from 2012 to 2017* for 37 Skill Categories

Panel A. Regression Analysis

	Dependent Variable: $\ln\left(\overline{duration}_{s,t}\right)$			
$\ln\left(\overline{application\_flows}_{s,t}\right)$	-0.079***	-0.039***	-0.032**	-0.030**
	(0.003)	(0.013)	(0.013)	(0.013)
$\ln\left(\overline{application\_flows}_{s,t-1}\right)$		-0.044***	-0.041**	-0.035**
		(0.013)	(0.017)	(0.017)
$\ln\left(\overline{application\_flows}_{s,t-2}\right)$			-0.011	-0.014
			(0.013)	(0.017)
$\ln\left(\overline{application\_flows}_{s,t-3}\right)$				-0.006
				(0.013)
Constant	2.382***	2.391***	2.394***	2.397***
	(0.006)	(0.006)	(0.006)	(0.006)
Observations	2,627	2,590	2,553	2,516
Adjusted R-squared	0.225	0.241	0.246	0.250
Within R-squared	0.136	0.157	0.160	0.164
Between R-squared	0.469	0.443	0.427	0.474

Panel B. Standard Deviation in Vacancy Posting Duration and Total Application Flows Within and Between Skills

	Mean Standard Deviation		**Median Standard Deviation**	
	ln(Avg. Duration)	ln(Avg. Applications)	ln(Avg. Duration)	ln(Avg. Applications)
Overall	2.225	1.993	2.232	1.864
Within Skill	0.081	0.697	0.077	0.660
Between Skills	0.138	0.566	0.138	0.561

Notes: We exclude skill categories with less than 25 distinct active postings per calendar month. We also exclude any job postings with first active date on or after December 1, 2017. We group job postings into skill-time categories based on the first skill requirement mentioned in the job title and the month-year of their first active date. For each skill-time category, we calculate average monthly posting duration ($\overline{duration}_{s,t}$) as the ratio of the sum of total vacancy days (measured in 24-hour intervals) and the number of vacancy postings. Average application flows($\overline{application\_flow}_{s,t}$) is the sum of total applications received divided by the number of vacancy postings in the group. We then regress the natural log of mean posting duration on contemporaneous log mean application flow. Each column after column (1) progressively adds an additional monthly lag of log mean applications as a regressor.

We calculate the within-skill R-squared as the squared correlation of the time-demeaned natural log of mean vacancy duration and the time-demeaned natural log of mean applications per vacancy. We calculate the between R-squared as the squared correlation of the across-skill mean in the natural log of mean posting duration and the natural log of mean application flows for each month-year.

Figure 1. The Distribution of Completed Spell Durations, Standard Vacancy Postings, January 2012 to November 2017

Panel A: Direct Hire Companies

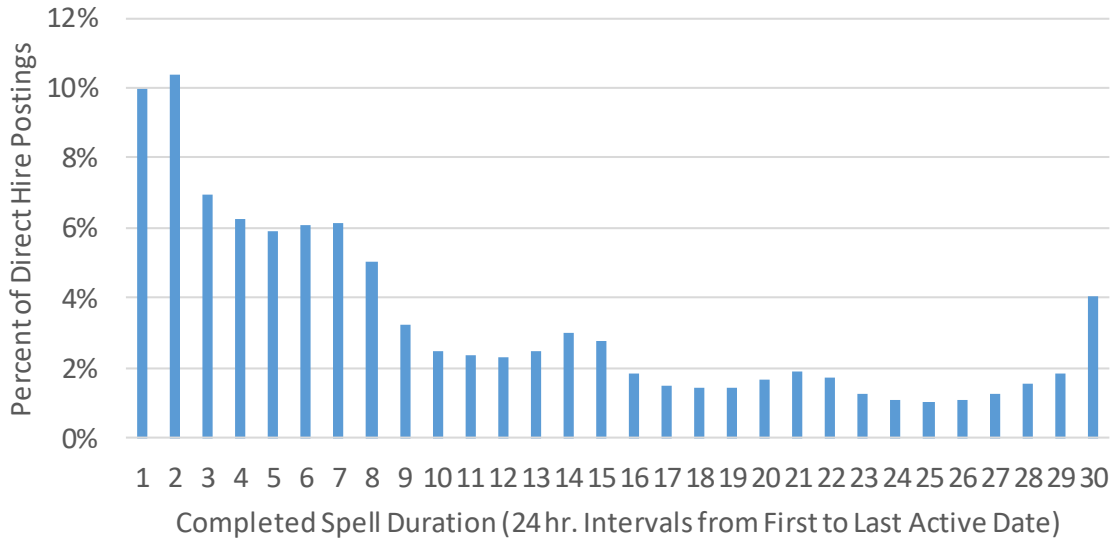

Panel B: Recruitment and Staffing Firms

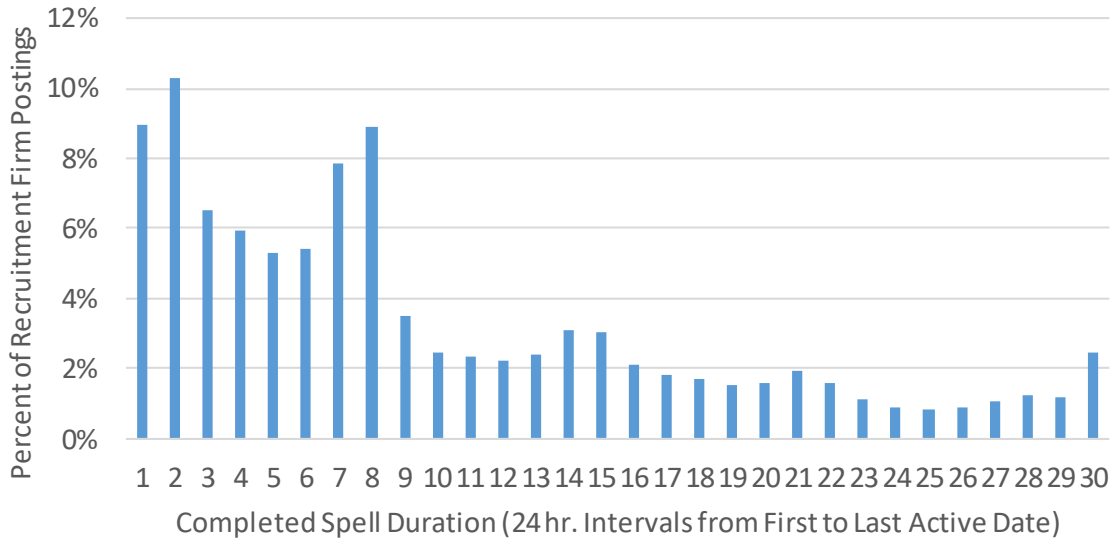

Notes: Jobs first posted on or after December 1, 2017 where removed to exclude incomplete spells.

Figure 2. The Distribution of Applications by Vacancy Posting Age, Standard Postings, January 2012 to December 2017

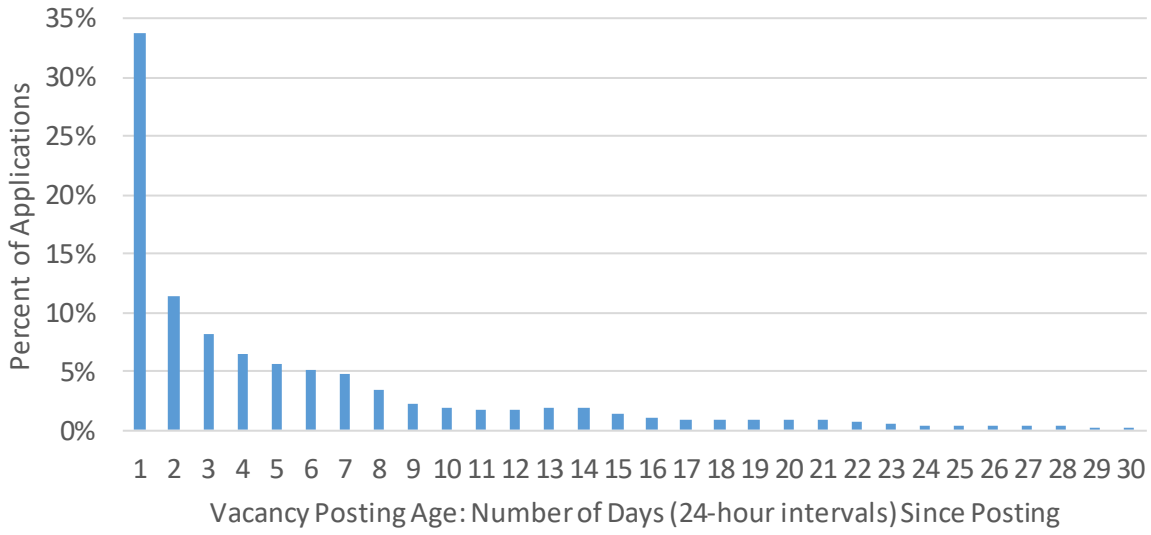

Figure 3. Mean Daily Applications Per Vacancy by Posting Age, Standard Postings, January 2012 to December 2017

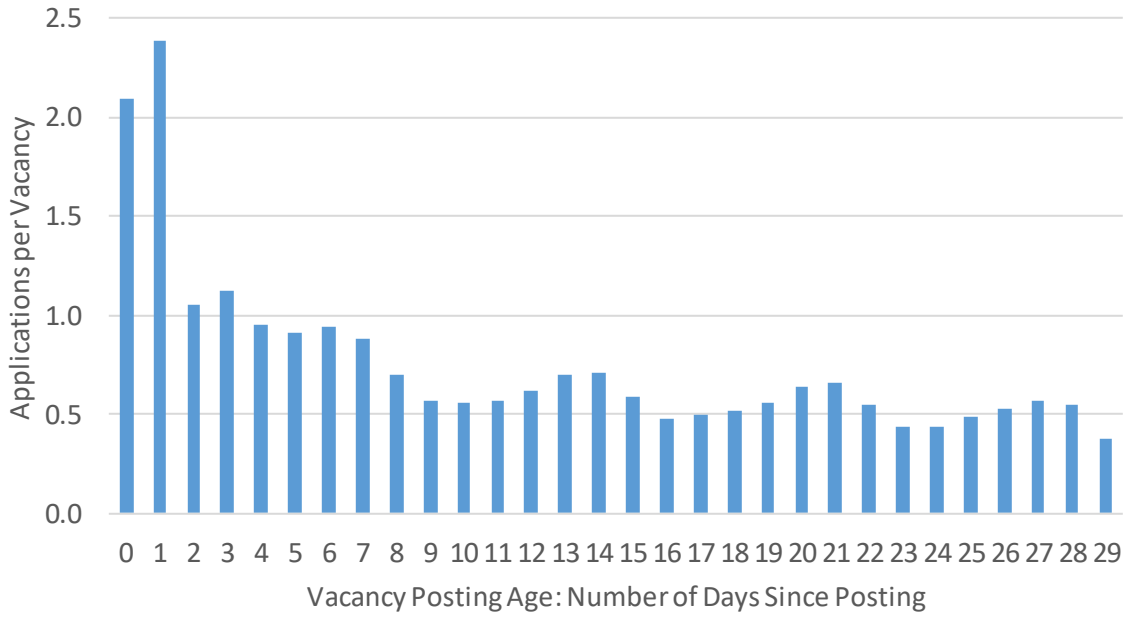

Notes: 0 in the x-axis indicates the day of first posting.

Figure 4. Frequency Distribution of Vacancies by Applications Received in First 14 Days Since Posting, Standard Postings, January 2012 – December 2017

Panel A: Direct Hire Clients

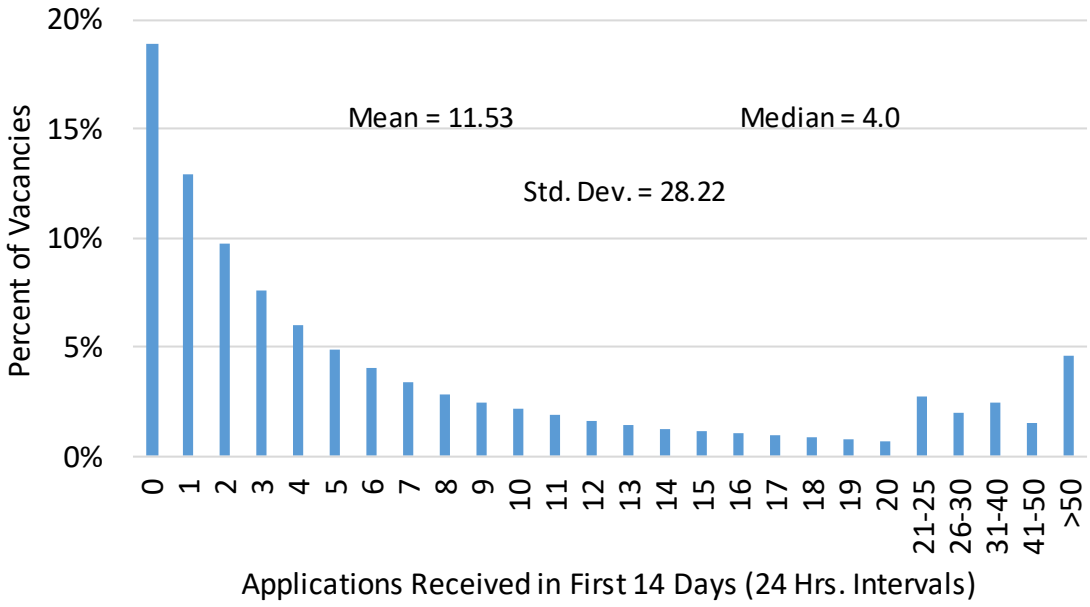

Panel B: Recruitment and Staffing Firms

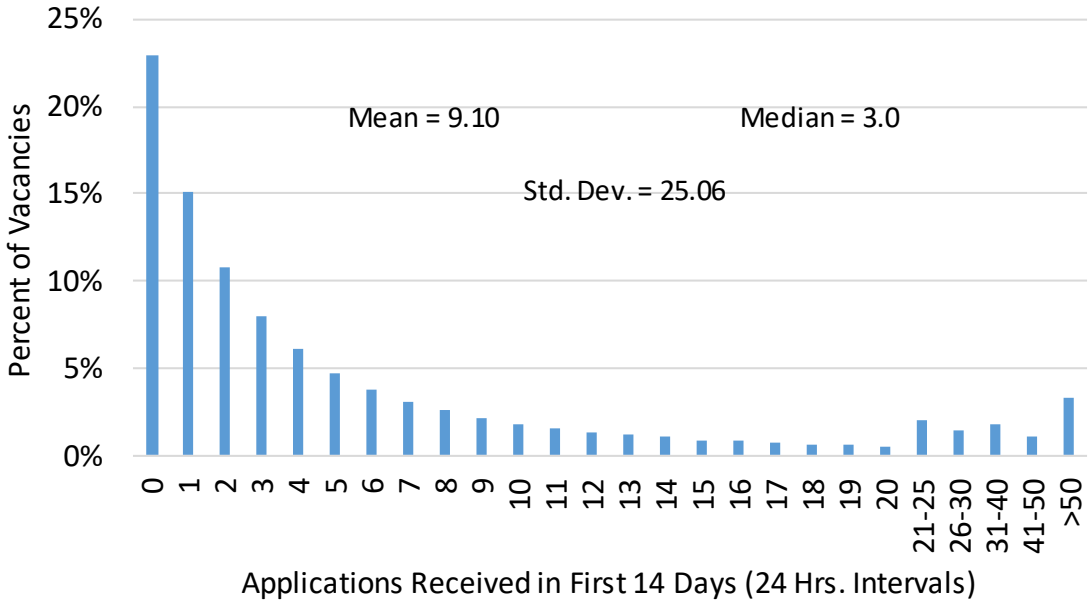

Figure 5. Mean Applications per Vacancy by Employer Size, January 2012 – December 2017

Panel A. Direct Hire, Standard Postings with Job Titles with at Least 100 Standard Postings, Equal Weights

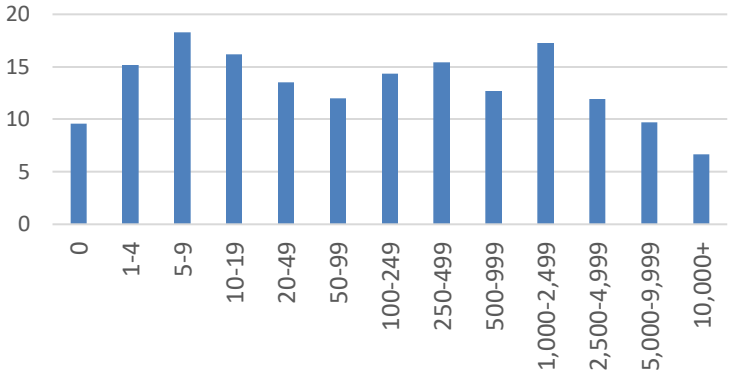

Panel B. Direct Hire, Standard Postings with Job Titles with at Least 100 Standard Postings, Weighted by Application Flows

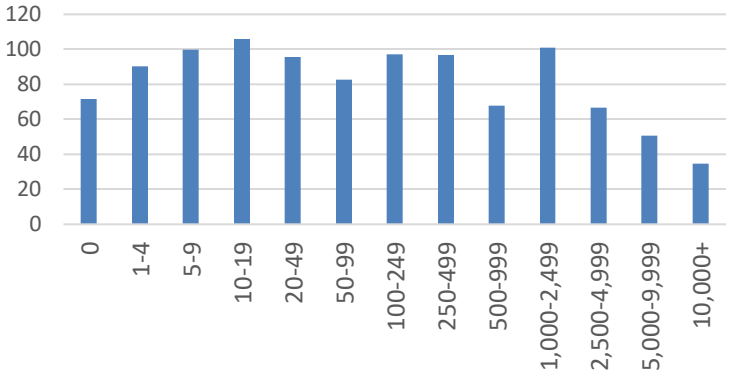

Panel C. Direct Hire, Standard Postings with Job Titles with at Least 100 Standard Postings Controlling for Job Title Composition, Equal Weights

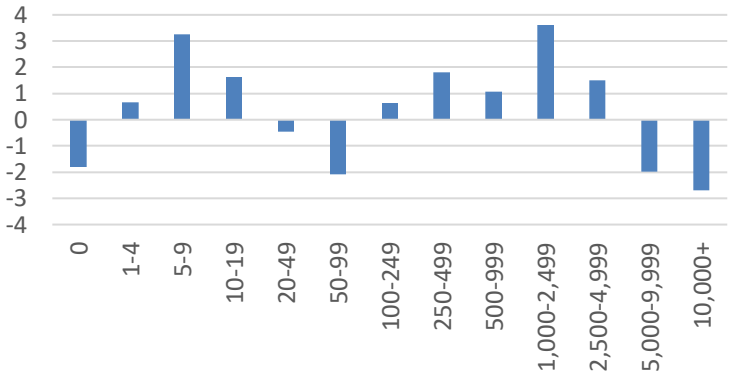

Panel D. Direct Hire, Standard Postings with Job Titles with at Least 100 Standard Postings Controlling for Job Title Composition, Weighted by Application Flows

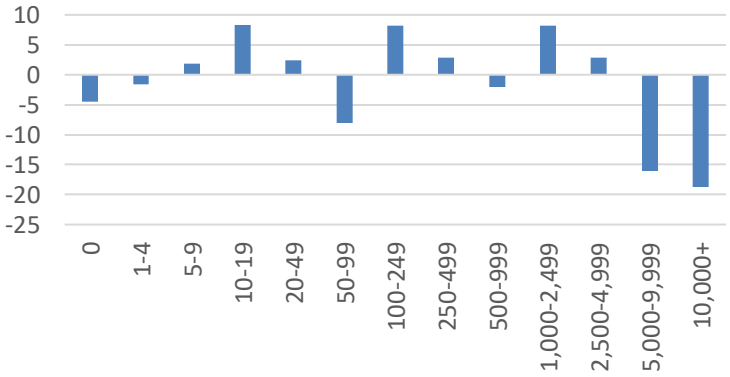

Note: X-axis shows employer size by number of employees in all panels. We obtain nearly identical results for Panels A and B if we consider all standard postings.

Figure 6. Applications Distribution. Moments implied by the ZINB model for selected values of the overdispersion and zero-inflation parameters compared to the DHI data

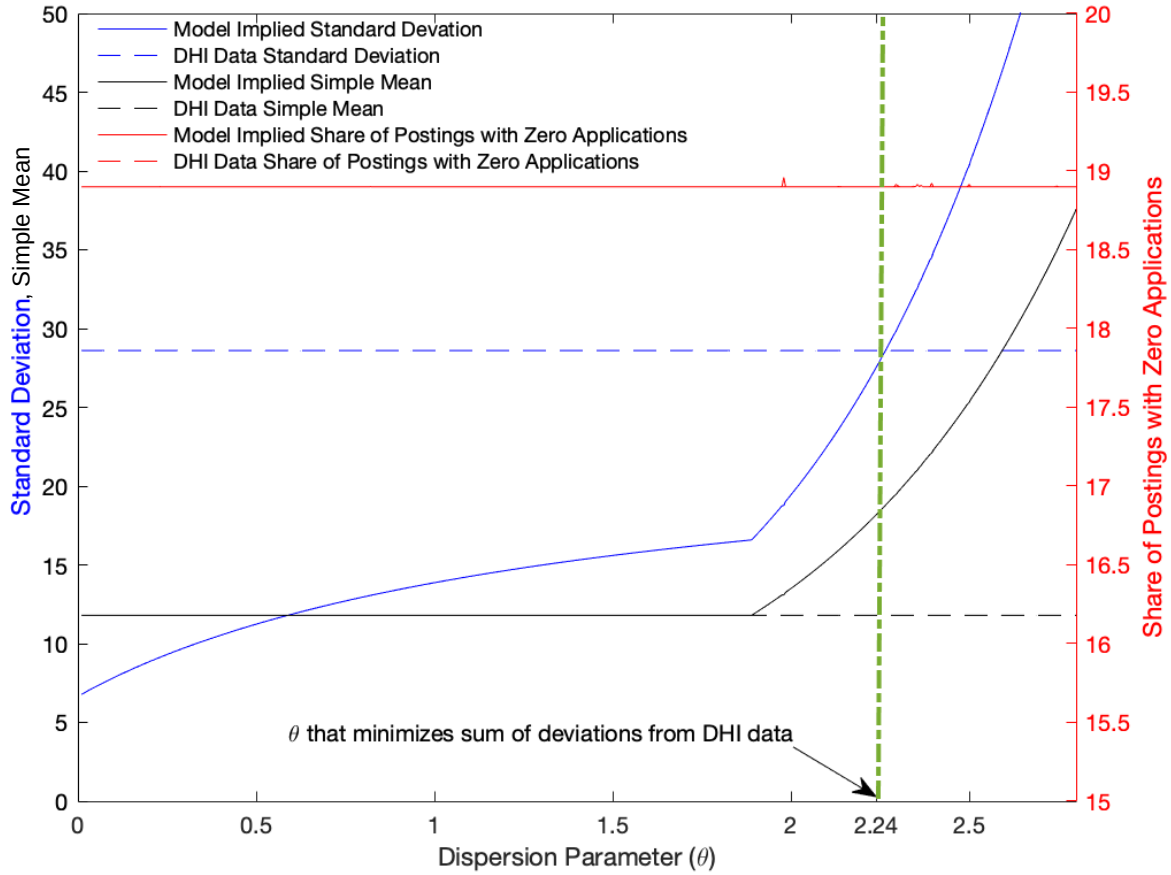

Notes: This figure shows the implied simple mean applications per vacancy, share of postings with no applications and standard deviation in applications by a zero-inflated negative binomial model as the overdispersion parameter increases. For each value of the overdispersion parameter, θ, we choose the other two parameters of the ZINB model (the mean of the negative binomial distribution, μ, and the zero-inflated probability, π) to minimize the sum of the absolute value of the deviation between the model implied and the empirical simple mean and the absolute value of the deviation between the model implied and the empirical share of postings with zero applications. As shown in Table 6, a model with an overdispersion parameter equal to 2.24 matches the standard deviation (26.8) and the share of postings with zero applications (18.9%) observed in the data but overshoots the mean applications per vacancy (18.8 vs. 11.8 in the data).

Figure 7. Stages of the Hiring Process and Relationship to Selected Data Sources

See main text.

Figure 8. Posting Duration and Realized Daily Application Flows, Conditional on Labor Market Tightness and Job Function Fixed Effects

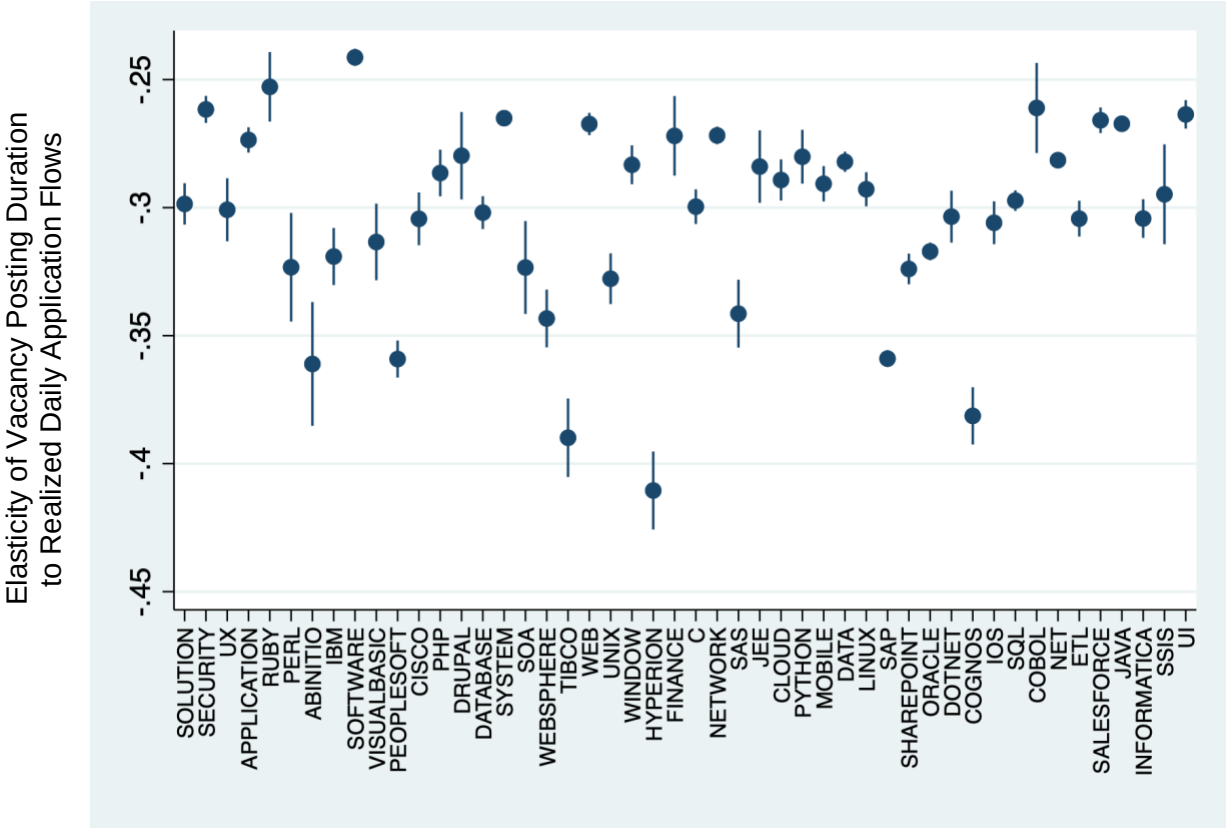

Notes: Each dot is the elasticity of posting duration to realized daily application flows, controlling for job function fixed effects, and indicator variable for whether the postings received any applications, and labor market tightness measured as mean applications per posting within skill-time cells. We order the skill categories in the x-axis from lowest to highest mean total application flows. This analysis excludes skill categories with fewer that 25 active postings in any calendar month.

Figure 9. Applicant Distribution by Search and Spell Duration

Panel A. Total Search Duration

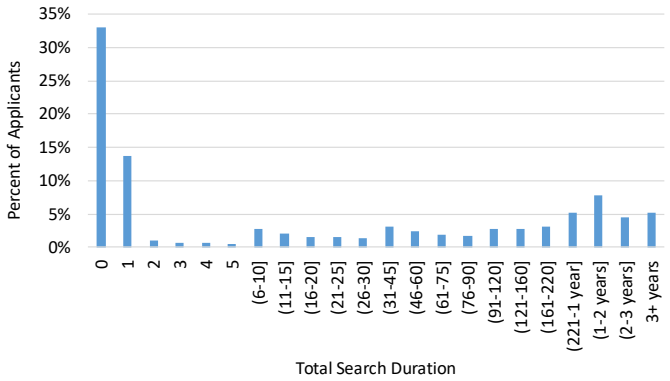

Panel B. Spell Duration

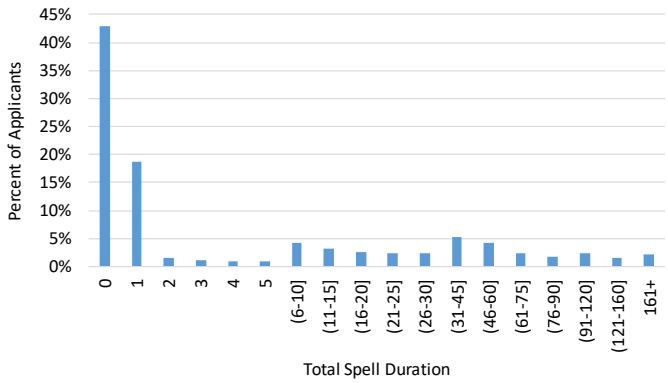

Notes: 0 in the x-axis refers to applicants that submit their last application on the same 24-hour interval as their first application. It also includes applicants that only submit 1 application. Search duration is the total time elapsed from a job seeker's first to last application on the DHI platform. We define spells as a group of applications where the time elapsed between consecutive applications does not exceed 60 days. Spell duration is the time elapsed between the first and last application within a spell.

Figure 10: Mean Applications per Applicant by Search Spell Duration at the Time of Application

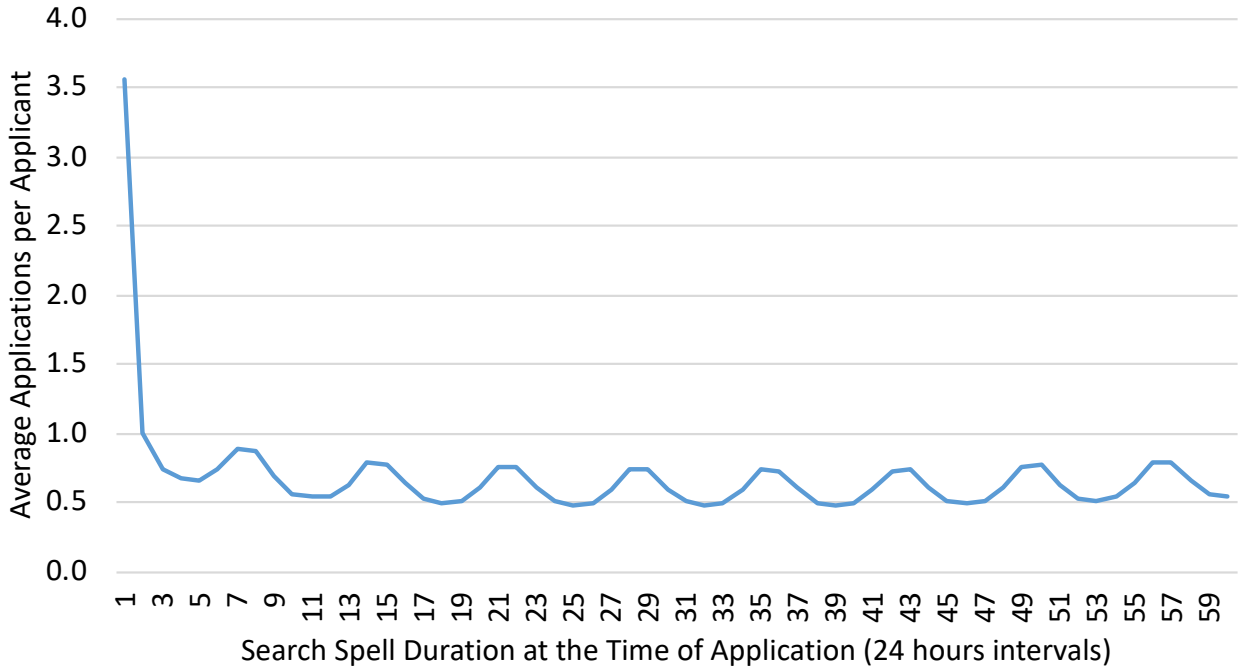

Figure 11: Application Volume Over Time, Standard Postings, January 2012 – December 2017

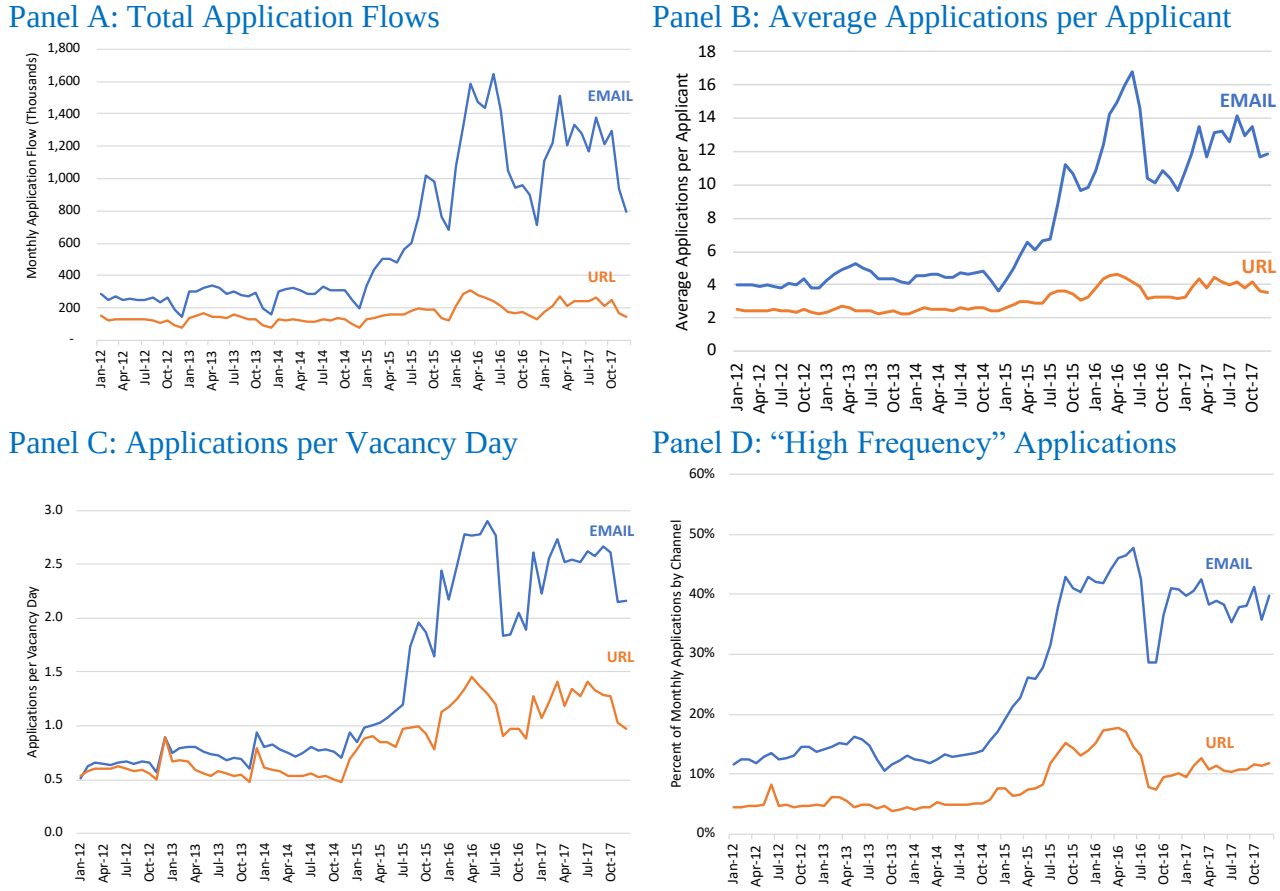

Notes: Panel A reports total monthly application flows to standard job postings by application channel. Panel B divides monthly application flows to standard job postings by the number of distinct job seekers (applicant IDs) that submitted at least one application via the corresponding channel (EMAIL or URL) to a standard posting in the month. In Panel C, we first group postings by the month and year of their first active date, and their application channel. We then divide the sum of total applications by each group of postings by the sum of their total active days. Total active days is measured as the number of calendar days when a posting was visible to job seekers. This figure excludes job postings that did not receive any applications. Panel D presents the sum of "high-frequency" applications received by standard job as a share of total monthly application flows to standard job postings by application channel. An application is a high-frequency application if the same job seeker (applicant ID) applied to a different job through the same channel (EMAIL or URL) within the previous 60 seconds.

Figure 12: Applications by Employer Size, Direct Hire Standard Postings, January 2012 – December 2017

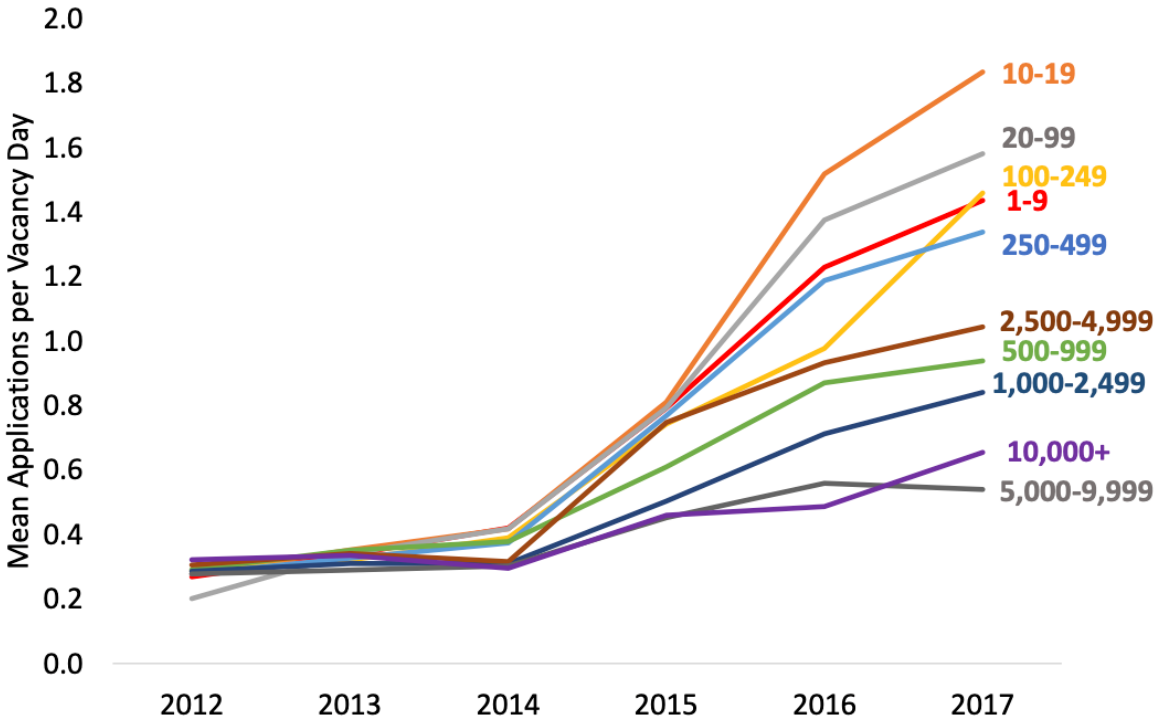

Notes: The graphs shows annual mean applications per vacancy day by employer size. The value for applications includes both email and URL applications. Postings standard job postings by Direct Hire clients, including those that received no applications. We exclude postings with first active date on or after December 1, 2017 and their applications.

Appendix A: Job Titles, Functional Groupings and Skill Clusters

Table A.1 lists the 30 most frequently posted job titles in the DHI Database, separately for Direct Hire Clients and for Recruiting & Staffing Firms. A full list of all 117,146 job titles with at least 250 distinct Job IDs (summed over both client types) is available as an electronic file at Job Posting Distribution by Job Title.

Table A.1. Most Frequently Posted Job Titles in the DHI Database

Direct Hire Clients		***Recruitment & Staffing Firms***	
Job Title	Job ID Count	Job Title	Job ID Count
DEVELOPER	88,510	DEVELOPER	223,713
ENGINEER	80,849	PROJECT MANAGER	183,936
MANAGER	62,407	ENGINEER	161,825
JAVA DEVELOPER	62,385	HELP / SUPPORT	161,614
PROJECT MANAGER	60,295	JAVA DEVELOPER	152,402
SOFTWARE ENGINEER	59,865	BUSINESS ANALYST	150,495
HELP / SUPPORT	51,497	ANALYST	119,302
ANALYST	50,694	MANAGER	93,206
BUSINESS ANALYST	50,380	NET DEVELOPER	92,036
CONSULTANT	45,866	CONSULTANT	80,902
ARCHITECT	35,922	SOFTWARE ENGINEER	72,508
LEAD	32,983	NETWORK ENGINEER	66,436
NET DEVELOPER	29,967	ARCHITECT	64,320
ADMINISTRATOR	28,628	ADMINISTRATOR	63,901
SENIOR SOFTWARE ENGINEER	26,833	WEB DEVELOPER	53,672
SYSTEM ENGINEER	26,608	TECHNICIAN	52,897
NETWORK ENGINEER	25,073	SYSTEM ADMINISTRATOR	49,516
SAP CONSULTANT	24,389	SENIOR JAVA DEVELOPER	49,241
SPECIALIST	22,855	SPECIALIST	48,845
SYSTEM ADMINISTRATOR	20,999	LEAD	48,167
SENIOR JAVA DEVELOPER	20,537	SYSTEM ENGINEER	41,374
SAP	20,325	SAP CONSULTANT	41,195
SENIOR ENGINEER	18,821	SQL DEVELOPER	36,885
WEB DEVELOPER	17,537	COORDINATOR	33,211
TECHNICIAN	16,318	DATA ANALYST	33,192
SALES	15,756	SENIOR PROJECT MANAGER	32,814
DIRECTOR	15,086	SENIOR DEVELOPER	32,040
SENIOR DEVELOPER	14,404	SAP	31,514
ORACLE DEVELOPER	13,081	C DEVELOPER	30,826
SOLUTION ARCHITECT	12,915	SYSTEM BUSINESS SYSTEMS ANALYST	30,660

Notes: "Job ID Count" equals the number of distinct Job IDs. The corresponding number of distinct Vacancy IDs for each title is larger due to the slicing operation described in Section II.B.

Table A.2 lists selected Job Function Categories, each of which aggregates over multiple job titles in the DHI Database.

Table A.2.1 Selected Job Function Categories in the DHI Database (Standard and Long Duration Postings)

Job Function Category	*Job ID Count*	*Weighted to Unweighted Mean Applications*		*Number of Job Titles*		*Average Weighted to Unweighted Mean Ratio Across Job Titles*	
		Per Vacancy Day	*Per Active Time*	*All*	*With at least 100 postings*	*Per Vacancy Day*	*Per Active Time*
Developer	1,603,009	4.8	3.2	515	213	3.8	2.5
Engineer	974,251	4.9	3.1	441	146	4.2	2.7
Administrator	520,918	3.2	2.2	347	116	2.7	1.9
Mechanical Engineer	9,916	2.4	1.3	49	6	2.3	1.7
Electrical Engineer	9,941	2.6	2.4	64	8	2.1	1.9
Business Analyst	295,523	2.6	1.8	279	56	2.6	1.8
Analyst	457,534	4.0	2.1	392	114	3.7	2.2
Help / Support Desk	331,791	2.6	1.7	341	73	2.5	1.6
Sales	60,196	4.2	3.3	188	29	4.1	3.1

Table A.2.2 Selected Job Function Categories in the DHI Database (Standard Postings Only Postings)

Job Function Category	*Job ID Count*	*Weighted to Unweighted Mean Applications*		*Number of Job Titles*		*Average Weighted to Unweighted Mean Ratio Across Job Titles*	
		Per Vacancy Day	*Per Active Time*	*All*	*With at least 100 postings*	*Per Vacancy Day*	*Per Active Time*
Developer	1,209,712	6.1	4.6	490	199	5.0	3.5
Engineer	649,279	6.9	4.5	414	131	5.8	3.8
Administrator	397,672	4.1	2.9	327	102	3.5	2.4
Mechanical Engineer	6,417	3.0	1.5	41	5	3.0	2.1
Electrical Engineer	6,311	3.1	2.8	49	6	2.4	2.1
Business Analyst	232,084	3.3	2.4	258	50	3.3	2.4
Analyst	335,447	5.2	2.7	368	103	4.8	2.8
Help / Support Desk	254,712	3.1	2.0	320	67	2.9	1.8
Sales	38,790	7.6	5.8	160	22	6.8	4.8

Similarly, Table A.3 aggregates job titles to obtain Skill Requirement Categories. To obtain the "SAP" Skill Requirement Category, for example, we first flag all Job IDs that contain "SAP" as part of the extended job title description. Job ID count refers to the sum of Job ID values in the Skill

Requirement Category. Number of job titles is the sum of distinct job titles that include the Skill Requirement. For a full list of all 54 Skill Requirement Categories see the file Job Posting Distribution by Skill Requirement Category.

Table A.3.1 Selected Job Skill Clusters in the DHI Database (Standard and Long Duration Postings)

Skill Requirement Category	*Job ID Count*	*Weighted to Unweighted Mean Applications*		*Number of Job Titles*		*Average Weighted to Unweighted Mean Ratio Across Job Titles*	
		Per Vacancy Day	*Per Active Time*	*All*	*With at least 100 postings*	*Per Vacancy Day*	*Per Active Time*
JAVA	419,895	4.7	3.3	212	57	4.2	2.7
SYSTEM	373,938	3.6	2.8	328	98	2.9	2.3
SOFTWARE	333,682	4.0	2.7	280	73	3.8	2.7
SAP	259,001	1.9	1.2	249	60	1.9	1.3
ORACLE	232,786	2.7	1.9	215	59	2.3	1.7
NETWORK	228,003	4.6	3.1	243	71	3.8	2.6
NET	214,321	4.7	2.6	199	43	4.3	2.4
DATA	187,084	2.9	2.1	289	66	2.5	1.8
APPLICATION	155,861	3.2	1.8	263	70	3.0	2.6
WEB	143,732	5.8	4.3	226	47	5.2	4.1
SECURITY	144,184	3.2	2.1	260	62	2.8	2.0
SQL	134,997	2.9	1.8	185	39	2.8	1.8
DATABASE	82,960	3.3	2.3	195	42	3.1	2.2
PEOPLESOFT	72,948	2.5	1.1	165	37	2.2	1.5
SHAREPOINT	71,826	2.8	2.0	178	35	2.6	1.9

Table A.3.1 Selected Job Skill Clusters in the DHI Database (Standard Postings Only)

Skill Requirement Category	*Job ID Count*	*Weighted to Unweighted Mean Applications*		*Number of Job Titles*		*Average Weighted to Unweighted Mean Ratio Across Job Titles*	
		Per Vacancy Day	*Per Active Time*	*All*	*With at least 100 postings*	*Per Vacancy Day*	*Per Active Time*
JAVA	312,933	6.2	4.9	198	50	5.4	3.8
SYSTEM	257,680	5.0	4.1	294	89	3.9	3.0
SOFTWARE	201,452	5.7	3.9	252	63	5.3	3.9
SAP	198,167	2.3	1.5	239	53	2.2	1.5
ORACLE	185,789	3.4	2.5	198	56	2.9	2.1
NETWORK	165,932	6.4	4.4	222	61	5.1	3.6
NET	163,889	6.1	3.4	190	38	5.5	3.2
DATA	138,621	3.7	2.7	268	58	3.1	2.3
APPLICATION	106,660	4.3	2.3	243	62	3.9	3.2
WEB	100,215	7.7	6.0	209	45	6.8	5.4

SECURITY	100,205	4.6	3.0	235	52	3.6	2.4
SQL	105,079	3.6	2.3	177	36	3.5	2.3
DATABASE	59,330	4.3	3.0	180	34	4.0	2.9
PEOPLESOFT	58,373	2.9	1.5	151	31	2.5	1.8
SHAREPOINT	56,215	3.4	2.5	165	31	3.2	2.3

Appendix B: The Job Seeker Experience on the Dice.com Platform

I. Introduction

In December 2014, DHI released a major update of its website. The main goal according to DHI was to help ensure that "the right type of applicant is able to find and apply for the right type of job". These changes resulted in a very significant increase in the number of applications per applicant received after 2014.

In order to extract accurate information about labor market slack and mismatch from DHI's data on vacancy postings and applicant flows we need to understand the changes that took place in 2014 and how they affected user experience and search outcomes. This document provides an overview of the changes to DHI's website in December 2014. It divides the changes into 3 major types:

1. Changes in the registration and application process,
2. Changes in search technology, and
3. Changes in matching technology.

For each of these categories, we discuss how they resulted in higher application volume below.

II. The 2014 Dice.com Release

The December 2014 release of DHI's job search engine (Dice.com) included several updates. These changes affected user experience, incentivized user's registration and profile completion, improved the relevance of search results and job suggestions, and enhanced overall website performance. This section provides further detail into these changes and associates them with the observed increase in application flows.

1. Changes in the Registration and Application Process

a. Registration

When an individual visits Dice.com, he is able to see and use the search engine to browse over 80,000 different job postings based on location, job title, company name, required skills, among other job characteristics. This search does not require being registered with DHI.

In December 2014, DHI implemented changes that not only made it easier to register but also increased the benefits from registration. Previously, users that started a registration process were asked to fill out a "skill card" that asked for basic skills information and submit a resume. Now,

registration only requires entering first and last name, email address, choosing an 8-character alphanumeric password and a CAPTCHA 3-digit number.[36]

Besides decreasing the cost of registration, DHI's website changes also increased incentives to register. Before 2015, several job postings included in their job description information such as an email address or application URL. This information allowed job seekers apply for jobs directly, without registering nor going directly through DHI's application channels. Therefore, these applications were not included in DHI's user statistics. Starting in January 2015, information that allowed direct applications was removed from job descriptions so that interested job seekers need to register in order to apply to jobs.

The total number of new registrations in Dice.com increased by 16% in 2015 relative to 2014. This is more than double the growth in registrations between 2013 and 2014 (7%). New registrations rose 62 percent from 46,607 in December to 75,733 in January. The percentage of unique visitors that chose to register rose from a monthly average of 2.5% in 2014 to 3.3% in 2015.

b. Application

Once applicants are registered they are prompted to complete a profile. Profile completion is optional and includes information on job seekers education, skills, desired position job title, work experience, among other relevant characteristics. Registered individuals can choose which, if any, of these fields to provide. They can also upload a resume and cover letter. The changes in the website in December 2014 make adding this information easier for individuals by adding drop down menus with options and a template with clear, separate fields for each of the relevant information that employers find useful (see Figure 1 below).

Figure B1: Dice.com new template for adding resume information

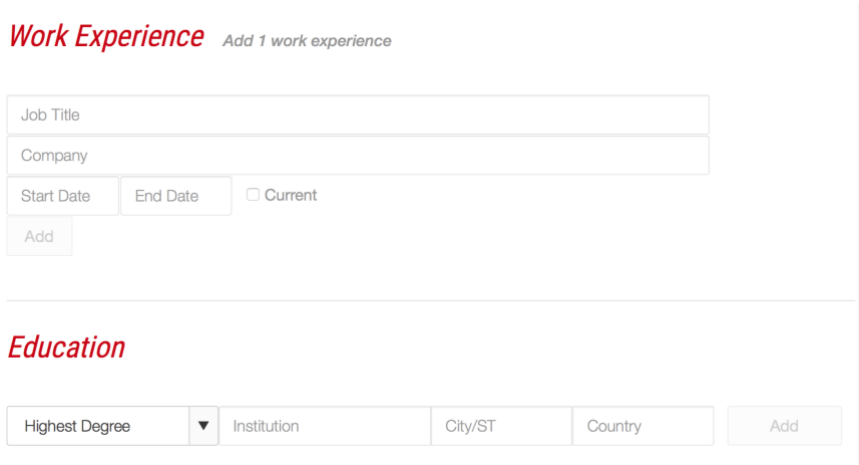

Completing a profile also makes it easier to apply to some jobs. Once the job seeker has added his resume to the website, applying for jobs that allow for direct applications through the DHI website[37] ("Email Applications") only requires selecting a few (at most 3) options from drop down

[36] The CAPTCHA used to be a question users had to answer.
[37] Job postings that allow applications directly through Dice.com are referred to as "Email Postings" in the DHI Database. See Davis and Samaniego (2016) for further detail on the DHI Database.

menus. Figure 4 below shows the application form for postings with email applications. The name fields are automatically field using the job seekers' profile information. The applicant then uploads a resume and can choose to also upload or write a cover letter. DHI provides the template cover letter shown below. Job seekers that are registered on the DHI platform can complete the application process in 10 seconds or less.

Figure B2: Email Applications Format

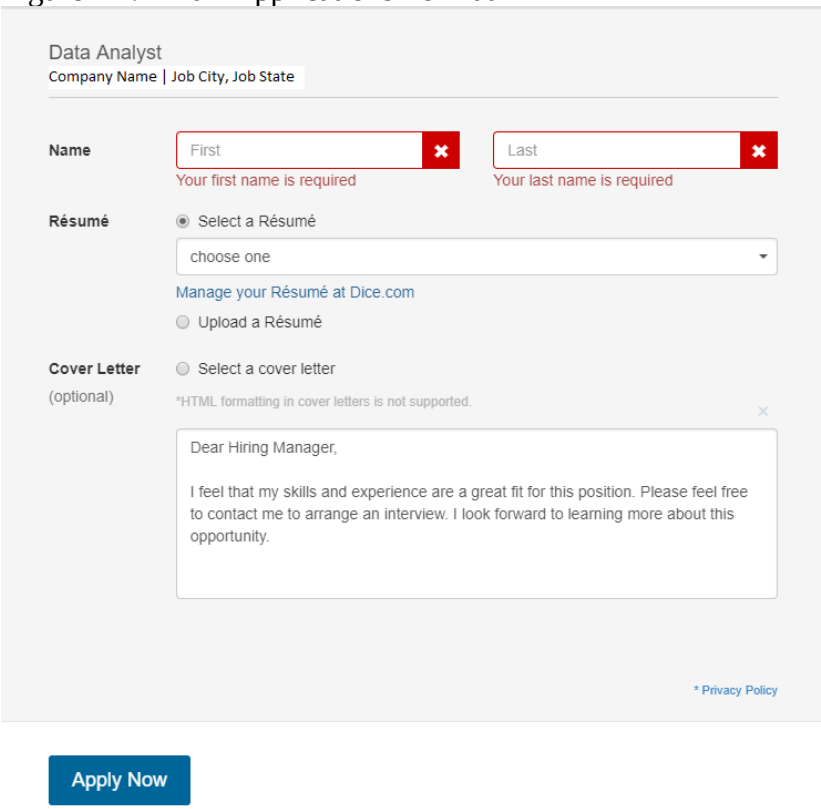

2. *Changes in Search Technology*

The changes implemented in December 2014 also improved the search process. Dice.com switched from Endeca, an older search engine, to SoIr. According to DHI, SoIr features better search functionality and configurability. In terms of user experience, this meant that job seekers were better able to tailor their search and find jobs with characteristics that they find desirable.

The new search engine allows job seekers to filter job postings based on the location of employment (either exact location or by mile range), company, job title, skill requirements, telecommuting options, among other keywords that limits the search. Job postings that fit the desired characteristics are then presented to the job seeker in order of relevance, that is, those jobs that are closest to the selected specifications appear first on the list.[38]

[38] It is important to note that when only the location is chosen as a filter, search results are ordered based on the posting date with the most recent postings appearing first.

Once a job seeker selects a job posting from the list of search results to look at the posting's detailed information, a list of similar jobs appears on the right side of the screen. This allows the worker to find several job postings with similar characteristics without requiring additional search efforts.

Figure B3: Similar Job Postings Suggested After Search

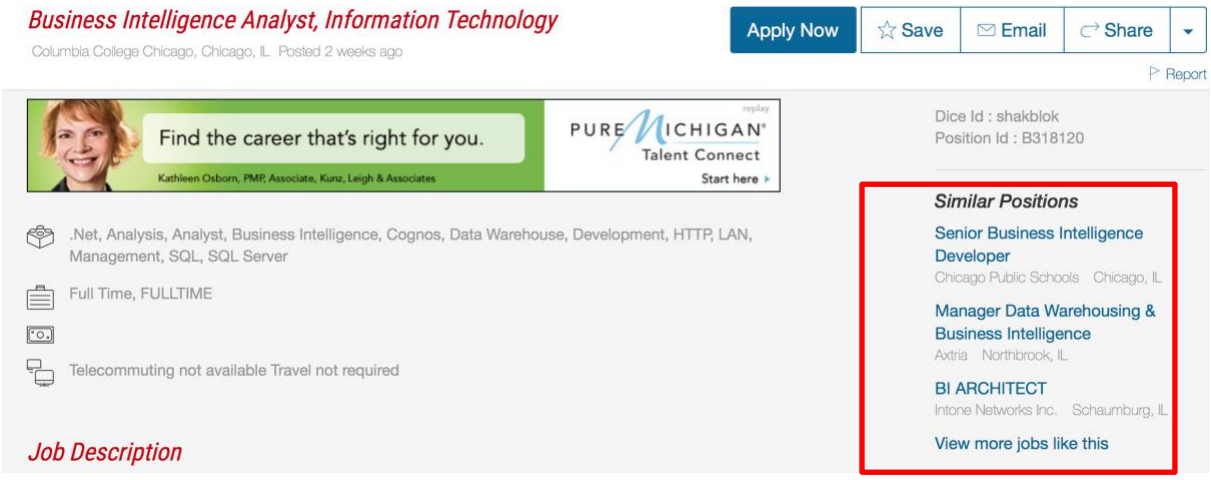

3. Changes in Matching Technology

Updates to Dice.com also allowed for better matching between job candidates and job postings that fit their skill sets. Individuals who complete their profile, can allow the information they included to be accessible to potential employers. Employers can browse the profile database and, if interested in a particular candidate, can express their interest in the candidates' application. Through this "searchable profile" functionality, employers can better assess the suitability of a candidate and potentially narrow their search efforts to those pre-screened job seekers. Moreover, this improved access to information for recruitment efforts also makes search more efficient for job seekers. Once a job seeker receives notice of interest for a particular job posting, he can also immediately see other similar postings from other employers (as shown in Figure 4 above).

These changes, couple with overall improvements in the website's performance and mobile access can help explain the significant increase in user registration and applicant flow observed after 2014. By making registration easier, more attractive and limiting the ability of non-registered users to apply through non-DHI managed channels, the updates to Dice.com increased the applicant flow included in the DHI Database. Moreover, by making search more efficient and applying easier, the website's updates increased the number of applications per applicant.

Appendix C: Additional Descriptive Statistics for Standard Postings

Table C.1 provides information about the distribution of completed spell duration by employer type and size.

Table C.1 The Distribution of Completed Posting Durations by Employer Type and Size, All Standard Postings in Job Titles with at Least 100 Standard Postings

			Percentile				
	No. of Standard Postings	Mean	10	25	50	75	90
All Job Titles with at Least 100 Standard Postings	5,157,666	9.44	1.00	2.93	7.00	14.02	22.71
Employer Type (ownership)							
Privately Held Companies	4,744,376	9.35	1.00	2.82	6.92	14.03	22.56
Publicly Listed Companies	258,737	11.15	5.00	7.00	8.00	13.84	24.04
Government	6,153	12.99	2.97	6.99	12.02	18.18	26.70
Subsidiaries	50	7.37	0.74	3.04	5.99	10.49	14.00
Other, e.g., NGOs	24	14.55	3.99	5.23	12.12	23.22	28.33
Missing Employer Type	148,326	9.14	0.99	2.83	6.80	13.83	22.00
Employer Size							
0 Employees	974,965	9.66	1.01	3.01	7.00	15.00	21.09
1-4	486,311	9.18	0.99	2.73	6.71	13.92	22.13
5-9	258,564	8.07	0.95	2.01	5.78	11.98	20.81
10-19	319,851	7.95	0.90	1.76	5.68	12.00	20.79
20-49	531,849	8.60	1.00	2.67	6.07	12.96	20.99
50-99	496,501	8.50	0.99	2.18	6.01	12.94	20.97
100-249	522,907	9.21	1.00	2.83	6.77	14.00	21.96
250-499	337,619	9.77	1.00	2.88	6.89	14.94	24.00
500-999	200,730	12.20	1.12	4.14	9.29	19.58	28.13
1,000-2,499	283,179	8.60	0.83	1.83	6.00	13.01	23.08
2,500-4,999	60,618	14.16	1.99	6.00	13.01	22.07	28.99
5,000-9,999	119,737	15.20	2.27	6.77	14.00	24.75	29.54
10,000+	420,332	10.44	2.01	6.00	7.83	13.75	24.00
Missing Employer Size	144,503	9.13	0.99	2.83	6.80	13.82	22.00

Notes: Table entries report statistics on completed spell durations for standard vacancy postings in job titles with at least 100 standard postings from January 2012 to December 2017. We measure duration from initial posting date-time to final removal date-time in seconds and express the statistics in 24-hour intervals. Information about employer type and size is obtained from Dunn & Bradstreet, typically when the client opens a new account and may not be current. In constructing this table, we dropped observations with first posting date on or after December 1, 2017 to avoid the inclusion of incomplete spells.

Figure C.1 displays the frequency distribution of monthly applications per vacancy posting for long-duration postings that are active on the first and last day of the month. Recall that we sliced the raw long-duration postings into monthly segments, as discussed in Section II.B of the main text. For Figure C.1, we restrict attention to long-duration postings that are active on the first and last day of the month ensure the suitability of a monthly applications count.

Figure C.1. Frequency Distributions of Monthly Applications Per Posting for Full-Month Long-Duration Postings (January 2012 – December 2017)

Panel A: Direct Hire Clients

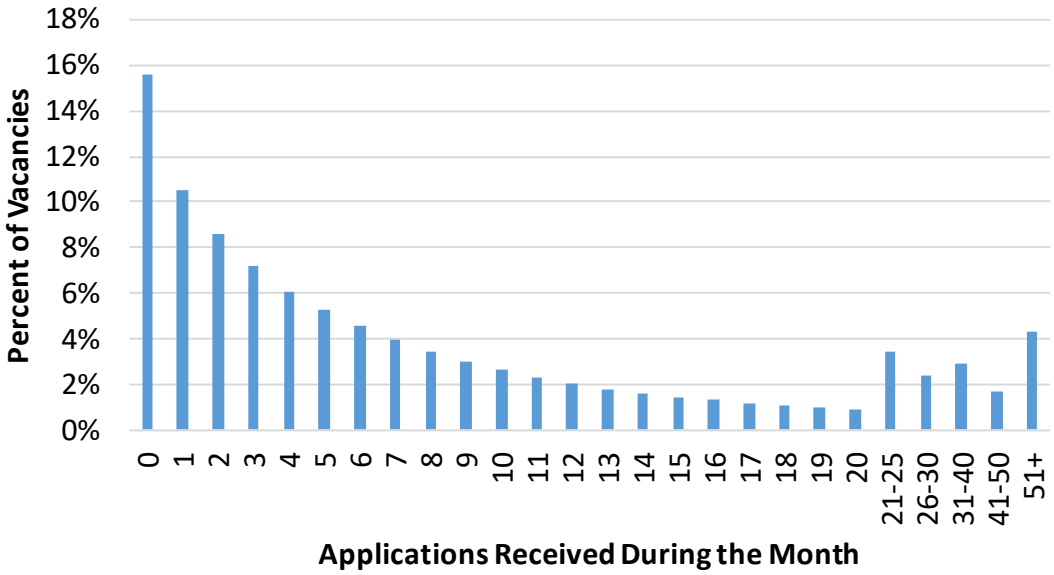

Panel B: Recruitment and Staffing Firms

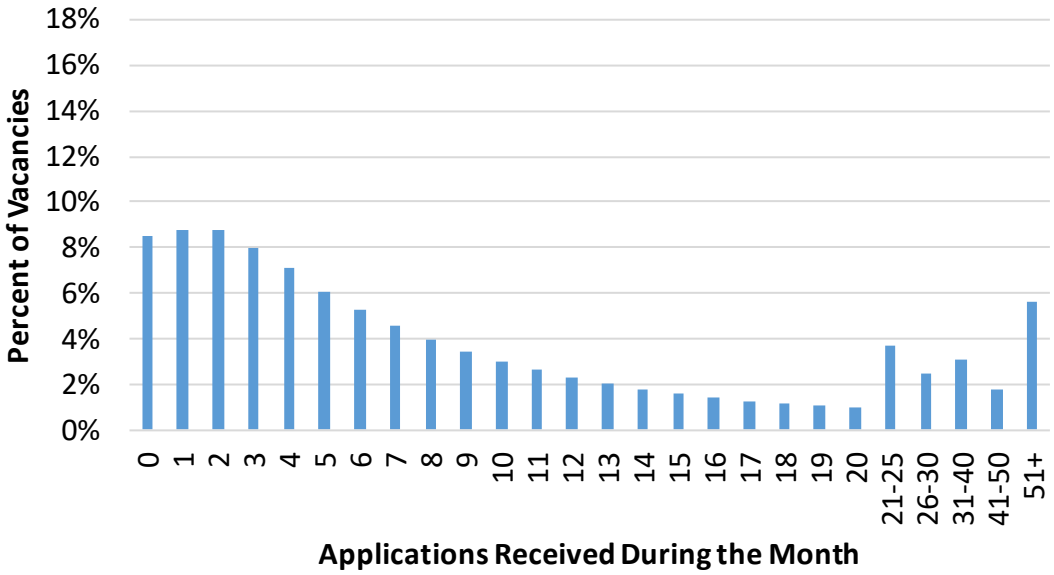

Figure C.2. Median Applications per Vacancy by Employer Size, January 2012 – December 2017

Note: X-axis shows employer size by number of employees in all panels.

Panel A. Direct Hire, All Standard Postings, Equal Weights

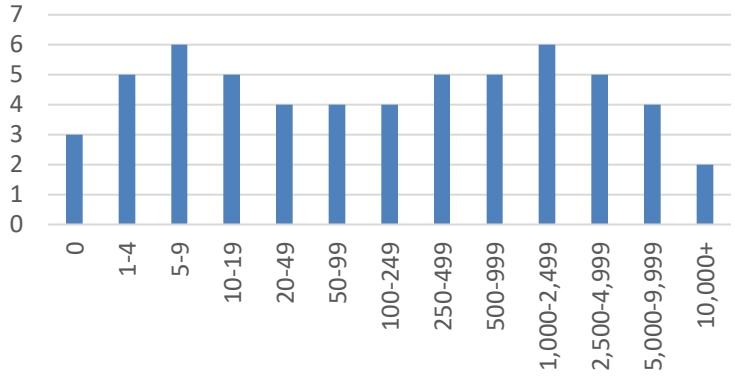

Panel B. Direct Hire, All Standard Postings, Weighted by Application Flows

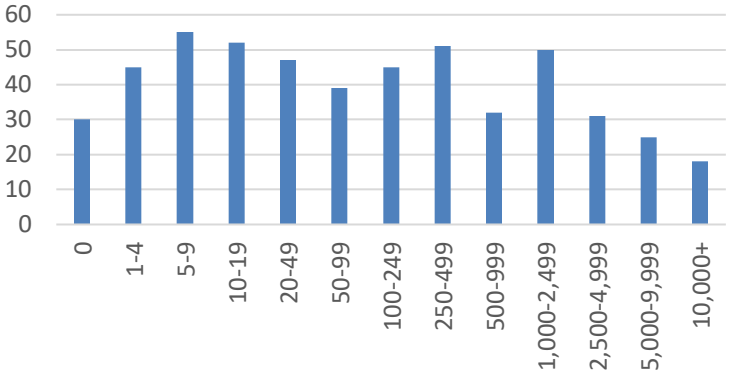

Panel C. Direct Hire, Standard Postings with Job Titles with at Least 100 Standard Postings, Equal Weights

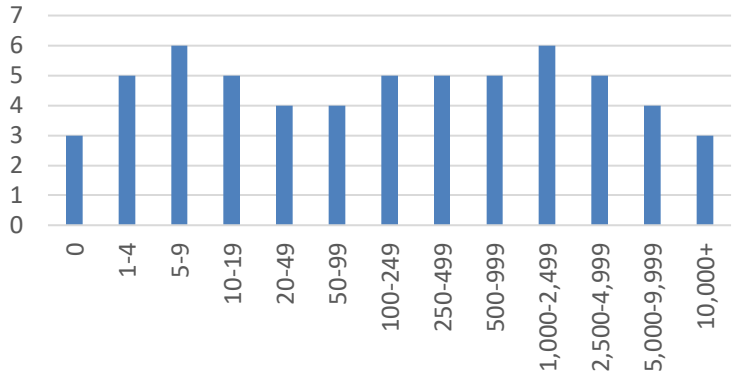

Panel D. Direct Hire, Standard Postings with Job Titles with at Least 100 Standard Postings, Weighted by Application Flows

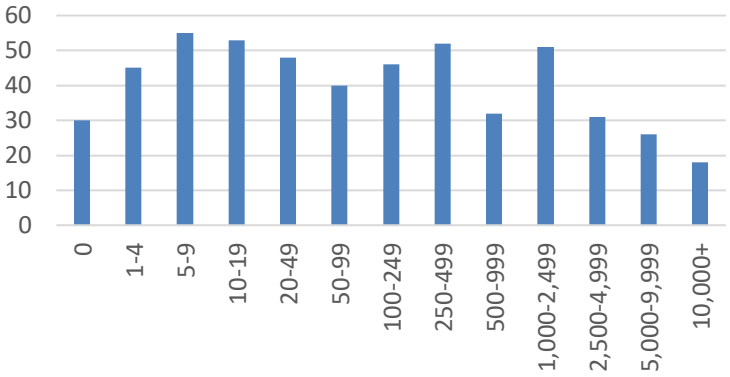

Panel E. Direct Hire, Standard Postings with Job Titles with at Least 100 Standard Postings Controlling for Job Title Composition, Equal Weights

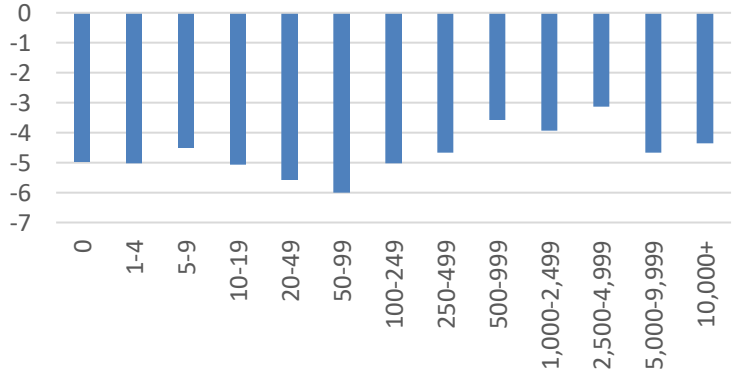

Panel F. Direct Hire, Standard Postings with Job Titles with at Least 100 Standard Postings Controlling for Job Title Composition, Weighted by Application Flows

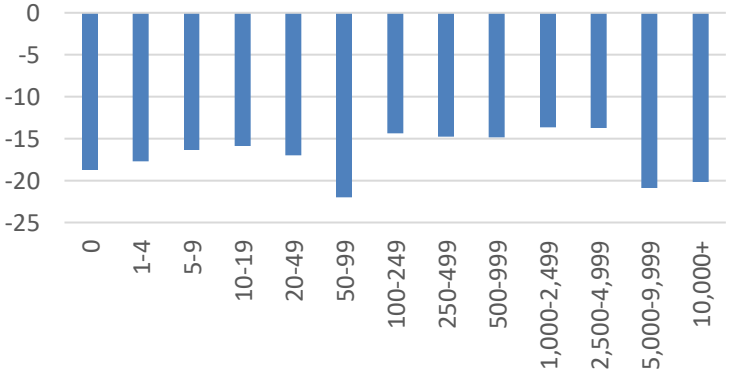

Figure C.3 displays the distribution of applications by posting age separately for Direct Hire Clients and Recruitment & Staffing Firms.

Figure C.3. The Distribution of Applications by Vacancy Posting Age, Standard Postings, January 2012 to December 2017

A. Direct Hire Clients

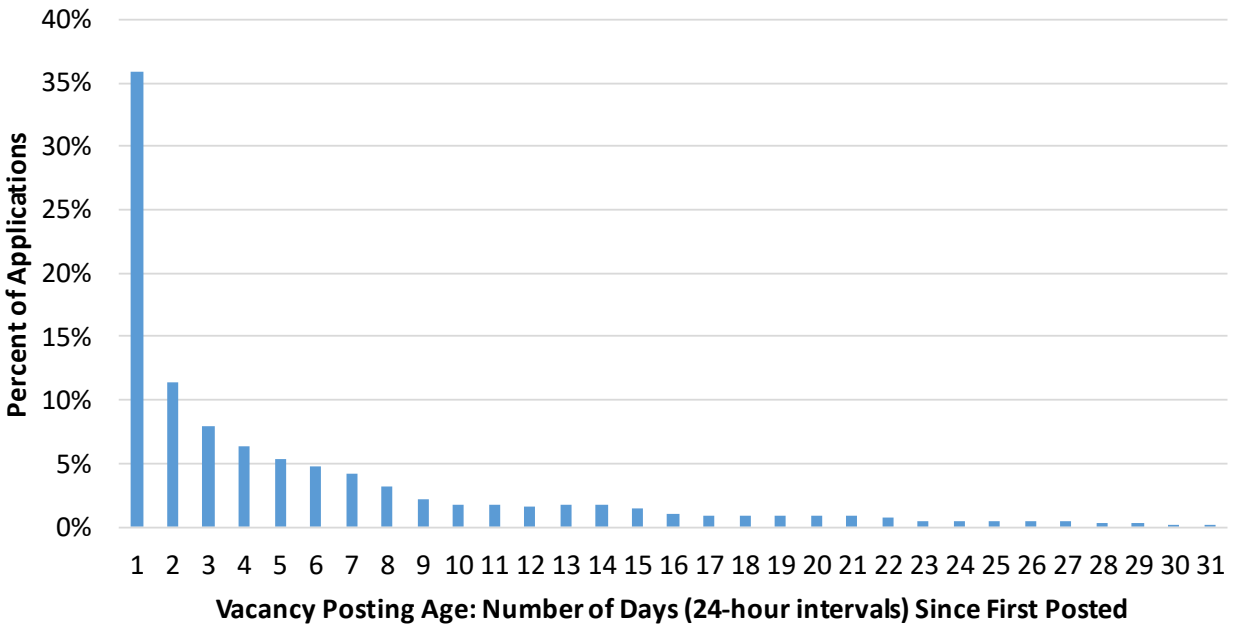

B. Recruiting and Staffing Firms

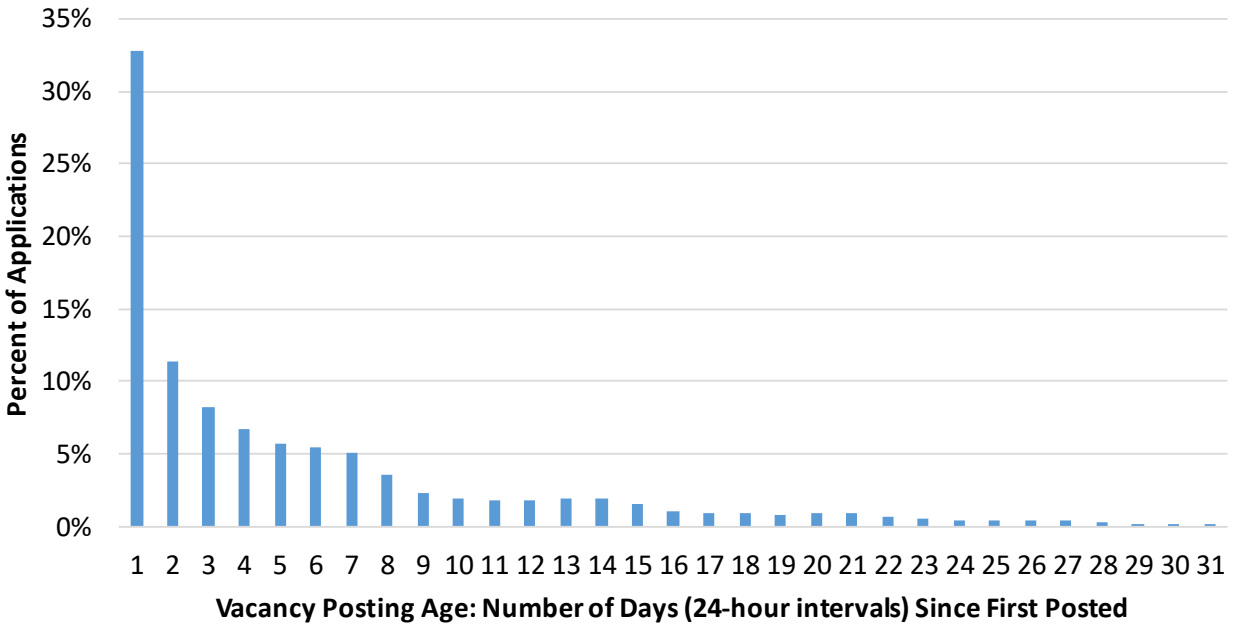

Appendix D: XXX

Our sample includes 125.3 million applications. 8.7% of these applications are duplicates, that is, they refer to an applicant ID applying to a job ID more than once. In our analyses, we keep these applications in the sample and treat them as distinct. We argue that it makes sense to keep these applications as they need not be treated as duplicates by employers. For 3rd party applications, multiple applications to the same job may refer to different applicants using the same intermediary to apply for jobs. For long-duration jobs, a job seeker applying more than once may be considered on a rolling-basis by the employer. The case for excluding duplicates is stronger in the case of URL applications where duplicates can arise if job seekers request to be redirected to the external application system more than once to complete a previously started application. Rather than applying different rules to include or exclude applications, we keep them all and present additional information about these applications below. Among other things, we show that the distribution of postings by total number of applications received is similar regardless of whether we exclude duplicate applications or not, indicating that duplicate applications are similarly distributed across job postings.

Table D.1. Share of Duplicate Applications, January 2012 – December 2017

Panel A: Standard Postings					
	Email		URL		Total
	Direct Hire	Recruitment Firm	Direct Hire	Recruitment Firm	
Total Applications (Millions)	17.4	31.3	4.4	9.2	62.3
Duplicates (% of previous row)	4.5	4.5	14.5	12.3	6.4

Panel B: Long Duration Postings					
	Email		URL		Total
	Direct Hire	Recruitment Firm	Direct Hire	Recruitment Firm	
Total Applications (Millions)	17.0	29.5	9.0	7.4	63.0
Duplicates (% of previous row)	9.0	8.8	17.2	16.7	11.0

Notes: The number of duplicate applications is equal to the difference between the total application flow and the sum of distinct applicant IDs. For each panel, the first row presents the total application flow. The second row shows duplicate applications as a share of total application flows by company type and application channel.

Table D.1 shows duplicate applications as a share of total applications by application channel (Email vs. URL) and company type (Direct Hire vs. Recruitment Firm) for standard and long duration postings. As expected, the share duplicate applications is higher for long duration postings, and for jobs that redirect to an external URL to collect applications. Job postings from Direct Hire and Recruitment Firms have similar shares of duplicate applications.

Appendix E: Out of Range Applications

We refer to "within-range" applications as those applications that have a time-stamp within the job posting's first and last active dates reported in the Activity File. Out of range applications are those that arrive outside of this time frame. Since job postings are only visible to applicants during active dates, when an application time stamp falls outside this range either the application's or the job posting's times are being misreported. We therefor exclude them (0.2% of all applications for US Direct Hire and Recruitment Firm companies; 0.6% of the full sample) from all our analyses.

Table E.1. Within and Out of Range Applications and Job Postings

Posting Type	Within-Range Application	% Total Applications	% of Total Postings*
Long Duration	No	0.2	0.6
Long Duration	Yes	93.1	85.6
Standard	No	0.0	0.2
Standard	Yes	6.7	13.6

Notes: This table presents the share of applications and postings affected by out of range applications. *The last column shows the share of job postings that either only received within-range applications or that received at least one out of range application.

Appendix F: Proof of Lemma Used in Section III.2

Lemma: Consider v vacancy postings, where v_n of them have $n = 0, 1, 2, \ldots, n^{max}$ applications. Let M and σ^2 denote the mean and variance of applications over postings, and let M^W denote the applications-weighted mean number of applications per posting. $M^W = M + (\sigma^2/M)$.

Proof: Let a be the total number of applications, and let a_n be the number at postings with n applications. The probability function of postings over the number of application is $f(n) = v_n/v$ for $n = 0, 1, 2, \ldots, n^{max}$. The probability function of applications over n is $g(n) = a_n/a = nv_n/a = nf(n)/M$, since $M = a/v$.

Using the relationship between the two probability functions, we can write the applications-weighted mean number of applications per posting as

$$M^W = \sum_n n\, g(n) = \left(\frac{1}{M}\right) \sum_n n^2 f(n) = \left(\frac{1}{M}\right)(M^2 + \sigma^2) = M + (\sigma^2/M)\ .$$

Q.E.D.

Section III.2 considers a random assignment of applications to vacancy postings. In this case, the number of applications at a given posting is a random variable distributed according to a binomial distribution with a mean of $M = a/v$ and a variance of $\sigma^2 = (a/v)[1 - (1/v)]$. It follows immediately from the lemma that $M^W = (a/v) + 1 - (1/v)$.