ECONOMICS WORKING PAPERS

Application Flows

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

Economics Working Paper 24102

HOOVER INSTITUTION
434 GALVEZ MALL
STANFORD UNIVERSITY
STANFORD, CA 94305-6010
October 2, 2026

We build and analyze a new U.S. database that links 125 million applications to job vacancies and employer-side clients on Dice.com, an online platform for jobs and workers in software design, computer systems, engineering, financial analysis, management consulting, and other occupations that require technical skills. We find, first, that posting durations are quite short, often only two or three days, with a median of seven days. Second, unlike vacancy durations, posting durations do not lengthen with market tightness. Third, job seekers display a striking propensity to target new postings, with almost half of applications flowing to openings posted in the past 48 hours. Fourth, applications per posting are much too uneven to reflect random search, even within narrow market segments and job categories. Moreover, explicitly posted offer wages play no role in explaining the deviations from a random-search benchmark. Fifth, intermediaries play a huge role on both sides of the platform: Recruitment and staffing firms account for two-thirds of all postings and attract most of the applications. We relate these and other findings to theories of labor market search.

Keywords: worker applications, job vacancies, intermediaries, tightness, vacancy durations, non-sequential search, staffing firms, recruitment firms

JEL Codes: J6, M5

Acknowledgements: We thank Yuri Bykov, Rachel Ceccarelli, Courtney Chamberlain, Jennifer Milan and Elizabeth Schillo for extensive consultations regarding the Dice.com data, platform, and business model. We thank Fernando Alvarez, Jason Faberman, Pieter Gautier, Philipp Kircher, Marianna Kudlyak, Ioana Marinescu, Ayşegül Şahin, Rob Shimer, Ronald Wolthoff, anonymous referees and seminar and conference participants at Arizona State University, the 2018 IZA Conference on Matching Workers and Jobs Online in Bonn, the 2018 SAM Conference in Frankfurt, the NBER Summer Institute, the San Francisco Fed, Singapore Management University, Stanford, UC-Berkeley, University of Chicago, Upjohn Institute, and Yale for many helpful comments on earlier drafts. The authors were compensated by DHI Group, Inc. for developing the DHI Database.

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

* The Hoover Institution at Stanford University and the Stanford Institute for Economic Policy Research. StevenD5@stanford.edu
# University of Notre Dame, Keough School of Global Affairs. bsamanie@nd.edu

I. Introduction

We study the flow of worker applications to job vacancies. To do so, we build a new U.S. database that links 125 million applications to vacancy postings and employer-side clients on Dice.com, an online platform for jobs and workers in software development, computer systems, engineering, financial analysis, management consulting, and other occupations that require technical skills. We obtained the raw data from DHI Group, Inc., which owns and operates online platforms for hosting job postings and attracting applicants. We worked extensively with DHI staff to process the raw data and to understand the Dice.com business model, platform, and user base. Our research database covers 7.5 million vacancies posted on Dice.com from January 2012 to December 2017.

We use the database to uncover several new findings about employer and worker search, and we relate the findings to leading search theories. First, *posting* durations for single-position openings are typically short, often lasting only two or three days. The median duration is seven days, and the mean is 9.4 days. The mean *vacancy* duration for comparable jobs in the Job Openings and Labor Turnover Survey is more than four times as long. Thus, the "meeting" phase of the search process, during which employers solicit and accept applications, is much shorter than the "selection" phase, which entails screening and interviewing applicants, picking one for a job offer, extending an offer, negotiating terms, and waiting for a decision to accept or reject the offer.

Second, posting durations show little sensitivity to labor market tightness, whether measured by the ratio of job openings to job seekers or applications per posting. In contrast, previous research shows that vacancy durations lengthen with market tightness,[1] confirming a central prediction of search models in the mold of Pissarides (1985, 2000) and Mortensen and Pissarides (1994) – hereafter, MP models. The implication is that screening, interviewing, selection and negotiation account for the cyclicality of vacancy durations. MP models, however, typically treat the meeting phase of the search process as coterminous with the vacancy spell, and greater tightness lengthens the overall vacancy spell by slowing the rate at which employers meet job seekers. Our evidence says the contact rate is not the source of that cyclicality: the stage during which employers gather applicants neither occupies much of the vacancy spell nor lengthens when markets tighten. MP models can still serve as a useful reduced-form characterization of the hiring process.

[1] See Davis et al. (2012, 2013), Crane et al. (2016), Gavazza et al. (2018), Leduc and Liu (2020), Mongey and Violante (2025), and Mueller et al. (2024).

Third, job seekers display a striking propensity to target new vacancy postings: 45 percent of applications on Dice.com flow to vacancies posted in the past 48 hours and 60 percent go to those posted in the past 96 hours. Application arrival rates drop sharply as postings age. Taken in isolation, this finding supports the empirical relevance of stock-flow matching models, as set forth in Coles and Smith (1998), for example. The bunching of applications in very young postings also favors a non-sequential search strategy, where employers first gather a pool of applicants and then offer employment to the most preferred applicant in the pool. [2]

Leading theories of frictional unemployment, including MP models, presume instead that employers assess each applicant on arrival and offer employment to the first one who passes a reservation threshold. This presumption is hard to square with the pattern of a brief employer-side meeting phase and a much longer selection phase. Thus, we see our evidence on the brevity of posting durations relative to vacancy durations, and the heavy bunching of applications shortly after posting, as motivation for models that feature non-sequential employer search. Early theoretical work on optimal search by Morgan and Manning (1985), for example, shows that non-sequential search on the employer side encourages workers to search non-sequentially as well. When we turn to worker-side behavior on Dice.com, we find strong indications of non-sequential search. We explain why the distinction between sequential and non-sequential search matters in section IV.3.

Fourth, application flows on Dice.com are distributed across postings in an extremely uneven manner. The unevenness is much too great to be rationalized as the outcome of random search. One potential explanation is that job seekers self-sort across labor markets defined by job location, skill requirements, and employer characteristics. We find that this type of sorting plays a major role in driving the unevenness of application flows, although we cannot show it yields a full explanation. Another potential explanation is that employers advertise wages in their postings to influence the direction of application flows, an idea that animates much theoretical research on directed search. Wright et al. (2021) review this literature. A basic problem for this explanation is that 83 percent of the postings on Dice.com do not state an offer wage or wage interval. For the other 17 percent, we find that offer wages play essentially no role in directing application flows or rationalizing departures from a random allocation. Thus, at least for job vacancies on Dice.com, explicit wage posting is a non-factor in explaining the distribution of application flows. While this

[2] We develop this point in Section IV.1, drawing on insights from Gal et al. (1981), Morgan (1983), Morgan and Manning (1985), and van Ours and Ridder (1992).

finding challenges the view that application flows are directed by posted wages, it does not contradict directed search more broadly, which may operate through non-wage attributes of the posting or through the expectations that job seekers form about compensation based on other information about the job or the firm.

Fifth, we find that 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. That is, intermediaries dominate activity on both sides of the Dice.com platform, which is itself an intermediary. Turning to the broader economy, we present evidence that staffing firms account for a sizable and rising share of worker allocations in recent decades. We also provide several pieces of evidence that point to a growing role for firms that provide headhunting, talent sourcing, screening and other recruitment services for clients that hire employees on their own account. All told, the evidence highlights the growing importance of intermediaries that provide recruitment and staffing services. There are sound reasons to think these intermediaries affect the character and quality of matching and other labor market outcomes, as we discuss in Section IV.4, though importantly, our findings on short posting spells, application bunching at young postings, and tightness-insensitive posting durations are not specific to postings and applications that involve these intermediaries.

Several other studies use data on applications and job postings to analyze search behavior. In early work with Dutch data, van Ours and Ridder (1992, 1993) argue that the combination of falling applicant arrival rates and rising fill rates as vacancies age is incompatible with sequential search by employers. van Ommeren and Russo (2014) provide evidence against sequential search for employers that rely on paid advertising or employment agencies to recruit applicants. Marinescu and Rathelot (2018) use applications and vacancies to quantify the contribution of geographic mismatch to U.S. unemployment in 2012. Banfi and Villena-Roldan (2019) investigate how application flows respond to wage information in Chilean job postings, and Marinescu and Wolthoff (2020) consider how they respond to job title and wage information in two American cities. Faberman and Kudlyak (2019) investigate how online application frequency varies with worker search duration. Modestino et al. (2020) investigate how skill requirements in job postings vary with worker availability. Rabinovich et al. (2024) consider the employer's decision of whether to state the wage in job postings, and how that decision varies with labor market conditions. We remark on other papers and branches of the literature below.

The next section describes the Dice.com business model and DHI Database. Section III develops several empirical findings about application flows and vacancy postings. Section IV offers evidence on the extent to which our findings hold for U.S. labor markets more broadly. There, we also explain why the distinction between sequential and non-sequential search matters for wage outcomes, job creation incentives, worker sorting, match quality, and the optimal design of unemployment insurance. Section IV also relates our results to search theory, which leads to some further empirical investigations. Section V offers concluding remarks. Appendices provide more information about data processing and the DHI database, report additional empirical results, prove a useful analytic result, and offer additional remarks about related literature.

II. The Dice.com Platform and DHI Database

Our database links 125 million applications to 7.5 million job postings and 57,000 employer-side clients from January 2012 to December 2017. The raw data derive from Dice.com, a platform owned and operated by DHI Group, Inc. We worked closely with DHI staff to build and document the database. Before describing the database in detail, we provide background about the Dice.com platform and business model, which inform our treatment and analysis of the data.

1. Dice.com Revenue Sources and Pricing

Dice.com generates revenues from employer-side clients for vacancy postings, access to résumé banks, and other recruitment services. During our sample period, 98% of job vacancies on Dice.com were posted under "Subscription" contracts that grant clients a specified number of "job slots."[3] This type of contract lets the client freely allocate postings to a given slot, provided the number of postings visible to job seekers at a point in time does not exceed the number of slots. The contract price varies with the number of slots and ancillary services. For example, DHI charges extra to scrape job postings from the client's website and repost them on Dice.com.

Given the pricing of slots, clients face an opportunity cost of keeping a given posting in active status, i.e., visible to job seekers. In particular, an active posting prevents the client from using the slot to post a different vacancy. Even when the cap on slots is nonbinding, the client has incentives to remove stale postings. For one thing, it is costly to respond to applicants. For another,

[3] DHI offered other vacancy posting options during our sample period, but they accounted for tiny shares of all postings. 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 in our sample period.

the employer-side client opens itself to reputational damage when it leaves stale postings in active status. This reputational concern is important according to DHI staff, partly because repeated interactions between job seekers and employer-side clients are common. In line with these remarks, we find that Dice.com posting durations are typically short, with a median completed spell duration of one week. Thus, we think our measured posting durations reflect the actual time intervals during which the client accepts applications. In contrast, stale postings are common on some prominent online job boards, leading to distinctive matching frictions and information externalities (Cheron and Decreuse, 2017, and Albrecht, Decreuse and Vroman, 2023).

2. Regulating the Applicant Pool on Dice.com

Third parties submit many applications on the Dice.com platform. As an example, consider a staffing firm with employees to lease. If the staffing firm identifies a suitable job posting for one of its employees, it can apply on the employee's behalf with the aim of leasing his or her labor services. Some employer-side clients want to receive third-party applications, and some do not. Dice.com lets employer-side clients specify whether to accept third-party applications for any given vacancy posting. It also offers other means to selectively filter applications, but these other means were unavailable or not widely used on Dice.com during our sample period.

DHI also takes other steps to regulate the applicant pool and enhance the value of Dice.com to employer-side clients. It relies on client complaints and other information to identify and deter "bad" behaviors and actors. An example of a bad behavior is a third-party application submitted to a posting that wants only first-party applications. An example of a bad actor is an individual or organization that submits many nuisance applications. DHI uses machine-learning methods to develop rules for blocking undesirable applications, including those from certain foreign locations, IP addresses and User IDs with a history of nuisance applications. After verifying that a candidate rule does not generate false positives, DHI implements it to block certain applications. Clients do not see blocked applications, and they are not part of our database.

3. The Job-Seeker Experience on Dice.com

Job seekers on Dice.com can register, create a profile, review vacancy postings, and submit applications free of charge. They can also freely access Dice.com career development tools and content about skill trends and salaries in local labor markets. Job seekers can browse and search postings by job title, job location, company name, skill requirements, and other job characteristics. Browsing and searching do not require registration, but a Dice.com visitor must register before

applying for a job. Registered users can also create a profile, decide whether to make the profile visible to others, and whether to upload a résumé. According to SEC filings, 81% of job seekers who post résumés 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 significant changes to the Dice.com platform in December 2014. These changes improved search functionality for job seekers, made it easier for job seekers to register on the platform, and streamlined the process for submitting applications to certain jobs. At the same time, Dice.com made it possible for employer-side clients to signal interest to registered job seekers who opt for a visible profile. These changes to platform functionality brought large increases in applications per posting on Dice.com, as analyzed in Davis and Samaniego de la Parra (2021), but they did not materially alter the empirical patterns we document in this study.

4. The DHI Database

The DHI Database identifies employer-side clients and records when they post and withdraw particular vacancies. The database includes information about each client's industry, size, organization type, and location. For each posting, we know the city of employment for the job on offer, the client's description of the job in the online posting, a unique Job ID that links to the employer's Account ID, and the date-time stamp for each application. We also know the exact number of seconds a posting was active (i.e., visible to job seekers) each day, and the number of views each posting receives on each day.[4] While Dice.com serves employers and job seekers in many industries, its 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. We restrict attention to jobs in the United States, which account for 99% of the vacancy postings in the database.

The database distinguishes two types of employer-side clients. "Direct Hire" clients, which post vacancies to hire their own employees, account for 82 percent of employer-side clients. Staffing and Recruitment firms account for the rest. Staffing firms hire mainly with the aim of leasing employees to other firms. Recruitment firms seek suitable job candidates for their clients to

[4] See Davis and Samaniego de la Parra (2019) for a complete description of the database, its file structure, variable definitions, and basic summary statistics for each variable.

consider, and they are more likely to use a single posting to recruit for multiple vacancies, jobs in more than one city, or jobs for multiple employers.[5]

When posting a vacancy, the employer-side client chooses between two application channels: In the "Email" channel, interested job seekers submit applications via the Dice.com platform. In the "URL" channel, job seekers who wish to apply for the position 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.[6] The client can select different application channels for different postings and can even change the application channel after posting, but that rarely happens.

As reported in Row (1) of Table 1, the DHI Database contains 7.5 million unique vacancy postings from January 2012 through December 2017, and these postings attracted 125 million applications.[7] Recruitment and Staffing firms account for 67 percent of postings and draw 62 percent of applications. Email applications (i.e., those submitted directly via Dice.com) account for 76 percent of all applications. Because we find similar patterns for Email and URL applications, we pool them in the ensuing analysis. Governments and NGOs account for less than one percent of Direct Hire postings. Accordingly, we interpret our results as pertaining to private sector behavior. Direct Hire postings are distributed widely by employer size (Table 2), and over 90 percent are posted by 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). Because listed firms are, on average, much larger and less volatile than privately held ones, the share of postings and gross hires accounted for by listed firms is smaller than its share of private sector employment.

By way of comparison, the BLS Job Openings and Labor Turnover Survey (JOLTS) reports 7.2 million end-of-month job vacancies over the same period in the Information sector, the closest counterpart to the job openings on Dice.com. As discussed in Davis et al. (2009, 2013), JOLTS data undercount vacancies for three reasons: time aggregation; some vacancies are out of scope for

[5] We formed these understandings through conversations with DHI managers and staff who work directly with DHI clients. While "staffing" and "recruitment" are distinct functions, the database does not distinguish between staffing firms and recruitment firms. In practice, the same firm often performs both functions, as we confirmed by reviewing the websites of several Recruitment and Staffing firms that operate on Dice.com.
[6] We know when an applicant clicks through to a particular external URL multiple times or applies multiple times via the Email channel. Appendix A details our treatment of these "repeat" applications. We have also confirmed that our findings are robust to simply excluding the repeat applications.
[7] About 0.2 percent of applications have a date-time stamp before the vacancy's initial posting or after its permanent withdrawal from the platform. We drop these out-of-range applications.

JOLTS or otherwise not captured by the survey; and the JOLTS sample frame underweights new and young employers, which account for a disproportionate share of vacancies. These observations and the fact that our database covers 7.5 million vacancy postings suggest that the Dice platform hosts a large share of all comparable job vacancies in the U.S. labor market.

Many vacancies in the DHI Database have short offline spells, whereby a given Job ID is posted, taken offline for hours or days, and then made visible again. These 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. We typically measure duration as elapsed time since initial posting, but results are similar when using cumulative time online net of offline spells.

Three-fourths of postings on Dice.com exhibit the following pattern: (1) The client posts a vacancy, (2) a large majority of applications arrive within the first week or two after posting, and (3) the client permanently removes the posting within 30 days after first posting. The data exhibit variations on this 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 vacancy posting for a single opening. (Conversations with DHI staff support this interpretation.) Other Job IDs stay online for many weeks or months, and applications flow in over time. Based on conversations with DHI staff and our examination of the data, the vast majority of these "long-duration" postings pertain to more than one job opening. They reflect clients with ongoing hiring needs for certain jobs, including Recruitment and Staffing firms that continually seek applicants for certain types of jobs. Hence, we focus on standard postings in the main text.

5. Classifications by Job Title, Job Function, and Skill Requirements

Marinescu and Wolthoff (2020) show the usefulness of job titles in classifying online 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 variation in the experience and education of applicants. Job titles are more useful in these respects than standard occupational classifications, because the titles contain more information about specialization, hierarchy (e.g., "staff accountant" versus "senior accountant"), and compensation. Hence, we use the text in the job-title field to construct detailed controls and to group vacancy postings by job titles, job functions, and skill requirements.

There are 1,983 job titles with at least 100 distinct postings (Job IDs) and 2,746 with at least 50. As seen in Table 3, the top 100 job titles account for 95 percent of the Job IDs in the DHI database and 96 percent of the applications. Appendix Table B.1 lists the most common titles. We also use the job-title text to group postings into Job Function and Skill categories. "Job Function" refers to our grouping of postings into 56 occupational categories such as "Programmer," "Developer," "Mechanical Engineer," "Consultant," and "Business Analyst." "Skills" refer to specific job requirements mentioned in the job-title text such as "C," "SQL," "Java," "User Interface," and "Big Data." Table B.2 reports summary statistics for selected Skill categories.

III. The Empirical Behavior of Postings, Applications, and Search

1. Posting Durations Are Short, Much Shorter than Vacancy Durations

Figure 1 shows the distribution of standard postings by completed spell duration, measured by time elapsed from initial posting to final removal.[8] Pooling data 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 bin covers durations from 24 to 48 hours. The duration distribution for postings by Recruitment & Staffing firms has a second mode at 8 days (168 to 192 hours), while the second most common bin for Direct Hires covers postings with durations of less than 24 hours. The mean posting duration is 9.8 days for Direct Hire clients, 9.3 days for Recruitment and Staffing firms, and 9.4 days overall.

The posting spell involves the solicitation of applicants. The vacancy spell also encompasses 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.[9] Davis et al. (2013) show how to use JOLTS data to calculate the mean vacancy duration. Applying their method to JOLTS data from January 2012 to December 2017 (and multiplying by (7/6) to convert from working days to calendar days), the mean vacancy duration is 40.2 days in the Information sector, the closest JOLTS

[8] 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.
[9] Employers may use other channels and platforms to attract applicants before, during or after the posting appears on Dice.com. Hence, our measure of the mean gap between posting and vacancy durations can reflect more than screening, selection, bargaining, and offer acceptance after the Dice.com posting ends. It can also reflect earlier or later efforts to attract applicants via other channels and platforms.

counterpart to the postings on Dice.com.[10] Thus, the mean vacancy duration is about four times as long as the mean posting duration.[11]

Table 4 presents information about posting durations by job function and applications volume, and Appendix Table C.1 provides analogous information by employer size and ownership type. Two results warrant particular attention. First, the median posting duration is a mere 7.0 days for Recruitment & Staffing firms and 6.9 days for Direct Hire clients. One quarter of all standard postings are active for 2.9 days or less.[12] That is, the "meeting" phase of the search and matching process is very short for a large share of postings. This characterization holds for all job types reported in Table 4, and it is broadly true of standard postings. Second, and somewhat to our surprise, completed spell durations tend to rise with application numbers. Of course, there is a mechanical effect cutting in this direction, as longer spells give more time for applications to arrive. Still, we had anticipated that employers would shorten posting durations in slack labor markets and lengthen them in tight ones. We return to this matter below.

2. Job Seekers Target New Vacancy Postings

Figure 2 displays the distribution of applications by posting age, defined as elapsed time since the posting first became active to the time of application. As the figure shows, job seekers exhibit a striking propensity to target new and recently posted vacancies: 45 percent of applications flow to vacancies posted within the previous 48 hours, and 60 percent go to those posted in the previous 96 hours. Older postings attract relatively few applications. Very similar patterns hold when we separately consider postings by Direct Hire clients and ones by Recruitment & Staffing firms (Appendix Figure C.1). For postings by Recruitment & Staffing firms, 44 (59) percent of applications arrive within 48 (96) hours. For posting by Direct Hire clients, the corresponding statistic is 47 (61) percent. Table 5 shows that a strong bunching of applications at freshly posted

[10] Mean vacancy durations are shorter in the broader non-farm economy during the early years of our sample period, but they draw closer together over time and converge by 2016.

[11] We also considered firm and posting characteristics to identify ex-ante criteria for excluding postings that represent continuous recruiting into positions with recurrent hiring needs. Specifically, we fit a highly interacted Lasso model to classify postings as "long duration" or "standard" based on the posting's job function, location, a set of firm size indicators, posted wage, and a set of industry indicators. Using the fitted Lasso model to classify postings based on their ex-ante characteristics yields a mean posting duration of 10.2 days, still only about one-quarter of the mean vacancy duration in JOLTS data for vacancy postings in the Information sector. Thus, the brevity of posting durations relative to vacancy durations continues to hold when we split postings on pre-determined variables.

[12] Including long duration postings increases the median posting duration from 7 to 11 days.

vacancies holds across quintiles defined by the volume of applications and across a heterogeneous set of job functions. The strong propensity for applications to flow to fresh vacancy postings is a ubiquitous feature of our data.

One reason fewer applications flow to older postings is because there are fewer old postings. In light of this fact, Figure 3 shows mean daily applications per posting by posting age. Postings receive, on average, 2.1 applications in their first day online and 2.4 applications on their second day. (A posting is often active for less than 24 hours on its first active day.) Afterwards, the application flow rate drops sharply to 1.0 per day and even fewer as postings age further.

The DHI Database records daily views for each posting, which lets us decompose mean daily applications per posting by age into views per posting and the conversion rate from views to applications. Views per posting peak at 22 on the first full day online and fall by roughly half by day nine. The conversion rate (mean applications per view) falls by about 45 percent from the day of posting to day two and by about 60 percent by day nine. See Appendix Figure C.3.

3. Many Postings Attract Few or No Applicants

Figure 4 displays the distribution of standard postings by number of applications received in the first 14 days online. For Direct Hire clients, 19 percent of 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. One-fifth of Dice.com postings attract no applicants, and one-third attract one or fewer applicants.[13]

It might seem surprising that many postings draw few applications. Three observations are helpful in this regard. First, most Dice.com postings specify demanding technical qualifications. Second, the job postings on Dice.com are concentrated in occupations with relatively rapid demand growth during our sample period, potentially outstripping the pace of skill adjustment on the supply side. For both reasons, we believe skill scarcities are more common for jobs on Dice.com than for the economy as a whole. Third, as we have discussed, DHI takes steps to block undesirable applications and regulate the application pool. These steps are part of DHI's efforts to provide high-quality applicant pools to employer-side clients.

4. Applications Are Distributed over Postings in a Highly Uneven Manner

[13] These results are robust to excluding postings that received zero views throughout the time that they are active on the platform.

While many postings attract few applicants, Figure 4 also reveals that 14 percent of standard postings by Direct Hire clients and 10 percent of those by Recruitment & Staffing firms attract 20 or more applicants within 14 days. More generally, Figure 4 shows enormous differences across postings in the volume of applications received. The highly uneven distribution of applications also holds when looking within groups of postings defined by job function, employer size, and other observables, and when restricting attention to job titles with many postings on Dice.com.

The extent of unevenness is much greater than can be explained by randomness in the flow of applications to postings. A few analytic observations help make this point. If *a* applications flow randomly to *v* postings, the number of applications at any given posting follows a binomial distribution with parameters *a* and $(1/v)$. The simple mean number of applications per posting is (a/v). The corresponding flow-weighted mean is $(a/v) + 1 - (1/v) \approx (a/v) + 1$ for large values of v. Thus, the flow-weighted mean number of applications per posting is only slightly greater than the simple mean under random search. More generally, let *M* and σ^2 denote the simple mean and variance of applications over postings, and let M^W denote the flow-weighted mean. Then $M^W = M + (\sigma^2/M)$, as proved in Appendix D. Thus, we can interpret the gap between the flow-weighted mean number of applications per posting and one plus the simple mean as a measure of distance from a random assignment of applications to postings.

As reported in Table 5, the simple mean of applications per standard posting on Dice.com is 11. The flow-weighted mean is 88, more than seven times greater than the value of 12 implied by random assignment. This result also holds within employer size classes (Appendix Figure C.2). Moreover, it continues to hold when we look within job functions, as illustrated in the bottom panel of Table 5. Among Electrical Engineers, for example, the simple mean is 3.7 applications per posting and the flow-weighted mean is 15.2. Among Business Analysts, the simple mean is 22.7 and the flow-weighted mean is 97.

It is also insightful to quantify the distance from randomness in another way. Given a random allocation, the expected fraction of postings that receive exactly *x* applications is $\frac{a!}{(a-x)!\,x!}\left(\frac{1}{v}\right)^x\left(1-\frac{1}{v}\right)^{a-x}$. For $v = 5.4$ million and $a = 59$ million, this formula implies an expected fraction of standard postings with no applications of 0.00002 percent. In the data, 20.4 percent of standard postings receive no applications. In other words, the observed share of postings with no applications is six orders of magnitude larger than the share predicted by a model with fully random

search. We return to the non-random allocation of application flows and the high share of vacancy postings with zero applications in Section IV. Among other things, we quantify the explanatory power of theories that stress the role of posted wages in directing job seekers to postings.

These results also say that the typical applicant faces many rivals for each sought-after job, even as employers face modest-sized applicant pools for most openings. The (unweighted) median number of applicants per posting on Dice.com is only four. In terms of economics, these patterns are consistent with two 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, a preferred location, or other desirable attributes. Second, that skill, geographic and other sources of mismatch are important phenomena that curtail the size of applicant pools for many vacancies and inhibit the matching of workers to job openings.

5. Intermediaries Play Major Roles on Both Sides of the Dice.com Platform

As a platform that facilitates matching between workers and jobs, Dice.com is a type of labor market intermediary. As it turns out, other intermediaries dominate activity on the Dice.com platform. Table 6 quantifies this point by presenting the joint distribution of applications over employer-side client types and worker-side application types. In the traditional conception of labor market matching, job seekers search on their own behalf by applying for jobs on offer by employers who recruit and hire on their own behalf. Remarkably, only 12% of applications on the Dice.com platform fit this traditional conception – these are the "1st-Party Applications" to "Direct Hire Clients." Among employers that hire on their own behalf (Direct Hire clients), more than 60 percent of their applications come from third parties, e.g., staffing firms. In addition, more than 60 percent of the applications by job seekers acting on their own behalf (1st-party applications) flow to openings posted by Recruitment and Staffing firms. Job postings by Recruitment and Staffing firms account for two-thirds of all postings on the platform. In short, intermediaries account for most of the activity on both sides of Dice.com, which is itself an intermediary.

6. Posting Durations Covary Strongly with Idiosyncratic Fluctuations in Applicant Numbers

How do posting durations vary with idiosyncratic fluctuations in applicant numbers? To address this question, we first estimate the following regression by least squares,

$$\ln\left(duration_j\right) = c + \sum_s \beta^s I\left[skill_j = s\right] \sinh^{-1}\left(daily\ apps_j\right) + s_j \times t_j + f_j + \epsilon_j, \qquad (1)$$

where j indexes postings, s denotes skill categories, t is the month the posting first became active, $s_j \times t_j$ are skill-by-month fixed effects, f_j are fixed effects for job functions, and ϵ_j is an error term. The inclusion of skill-by-month fixed effects controls for market-specific tightness and any other forces that vary over time at the level of skill groups. Specification (1) also controls for unobserved sources of variation across skill groups and job functions. The dependent variable is the natural log of the posting duration, measured as time elapsed from the first to last active date-time and expressed in days. The chief explanatory variable of interest is $\sinh^{-1}\left(daily\ apps_j\right)$, where $daily\ apps_j$ is the number of applications received in the first 14 days divided by 14 (regardless of posting duration). We interpret the β^s coefficients as skill-specific elasticities of posting durations with respect to conditional fluctuations in applicant numbers.[14] Our sample for (1) contains all standard postings in skill categories with at least 25 distinct postings in every month.

Figure 5 plots the estimated elasticities, which center at -0.41 and range from -0.28 to -0.59 across skill groups. To see what this means for the magnitude of the covariation, note that the cross-sectional standard deviation of $\text{asinh}^{-1}\left(daily\ apps_j\right)$ in our sample is 0.88 after deviating about skill-by-month and job function means. Thus, a two standard deviation increase in the number of applications to a particular posting (conditional on tightness, etc.) involves a reduction in the posting duration of 2(0.88)(0.41) = 72 log points. This is a large difference. It says that employer-posting durations shorten (lengthen) when applicant numbers are larger (smaller) relative to those received by other postings in the same skill-by-month and job function. Conditional on market tightness and job function, the fitted version of (1) implies that the variation in realized applications per posting accounts for 9% of the variation in posting durations.[15]

These estimates condition on job function, skill category, and skill-level variation in market tightness. They leave open the possibility that other characteristics of the job or the posting, unobserved to us but not to the job seeker, drive both applications and posting durations. Thus, in Appendix F we present results for a specification that includes posting-level fixed effects to absorb any fixed, unobserved attributes that influence posting durations. Those results are qualitatively and quantitatively similar, as we discuss in the appendix.

[14] The inverse hyperbolic sine transformation closely approximates the natural log transformation while accommodating zeros (e.g., Bellemare and Wichman, 2020). Using the natural logarithm of daily applications in the first 14 days plus 1 as the dependent variable yields similar results.
[15] We calculate the share of variation in posting durations due to variation in realized daily applications, net of market tightness and job function fixed effects, using the partial R-squared.

In unreported results, we re-estimate (1) after adding a control for whether the posting attracted at least one applicant. This extended version of (1) yields duration elasticities that are more than fifty percent greater (conditional on attracting at least one applicant), reinforcing the evidence that employers withdraw postings early when they attract enough applicants. Results for the expanded specification also show that (conditional) duration elasticities tend to be larger in magnitude for skill groups that attract fewer applications.

7. Posting Durations and Market Tightness

Previous research firmly establishes that *vacancy durations* lengthen with market tightness, as measured by the ratio of job openings to job seekers.[16] This empirical regularity confirms a central prediction of search models in the mold of Pissarides (1985, 2000) and Mortensen and Pissarides (1994). Previous research is largely silent about which aspects of search and matching account for cyclical movements in vacancy durations. Because tighter labor markets bring a slower pace of applicant arrivals in MP models, it is natural to hypothesize that the meeting phase of vacancy spells is longer in tighter markets and shorter in slack ones. To test this hypothesis, we treat posting spells as coterminous with the meeting phase of vacancy spells. Specifically, we investigate how *posting durations* vary with labor market slack.

To do so, we first compute the average posting duration by skill-month cell as total posting days in the cell (cumulated over active postings) divided by its number of postings. Second, we measure slack as the number of Dice.com job seekers that apply to one or more jobs in the cell divided by its number of postings. Third, we regress the cell-level mean posting duration on the cell-level slack measure. We control for skill fixed effects, because here we aim to uncover how posting durations covary with slack over time (not how they covary across skill categories).

Column (1) in Table 7 reports the results of this regression fit to monthly skill-category data from January 2012 to November 2017. The duration elasticity is negative, as hypothesized, but it is only -0.039 and precisely estimated. The time-series standard deviation of the log slack measure, averaged over skill categories, is 0.45. Thus, the fitted regression implies that a two standard deviation increase in log slack shrinks posting durations by 2(0.45)(0.039) = 3.5 log points, or one-third of a day when evaluated at the mean posting duration of 9.4 days. For perspective, U.S. vacancy durations rose from 18.6 days in July 2009 (the first month after the Great Recession) to

[16] See Davis, Faberman and Haltiwanger (2012, 2013), Crane et al. (2016), Gavazza, Mongey and Violante (2018), Leduc and Liu (2020), Mongey and Violante (2025), and Mueller et al. (2024).

39.3 days in September 2018.[17] Clearly then, slack effects on *posting durations* in the Dice.com data do not explain the sensitivity of *vacancy durations* to slack in much other research.

We subject this result to a variety of robustness checks. First, we obtain similarly small posting-duration elasticities when defining labor markets in terms of job functions, MSAs, or MSA-by-skill cells (Appendix Table C.2), and when re-estimating Table 7 separately for postings by Direct Hire clients and those by Recruitment & Staffing firms.[18] Second, changes to the Dice.com platform in December 2014 led to a strong rise in applications and applications per posting over the next several months (Davis and Samaniego de la Parra, 2021). Perhaps the platform design changes affected the ratio of job seekers to postings in ways that do not accurately reflect movements in slack. To address this concern, we refit the regression with controls for common time effects. Column (2) reveals that this specification yields a small positive posting-duration elasticity. Third, in unreported results, we added lagged values of slack to the regression model and again obtain similar results (for the sum of the coefficients on current and lagged slack values).

Finally, we consider an alternative slack measure motivated by models in which workers can submit multiple applications at the same time and employers collect a pool of applicants before interviewing some of them. Examples include Albrecht et al. (2006), Galenianos and Kircher (2009), Kircher (2009), Albrecht et al. (2020), and Cai et al. (2025). In the model of Albrecht et al. (2020), for example, the number of applications per posting is a fixed multiple of job seekers per posting, the value of which depends on the cost of an application.[19] If this property holds exactly in the data, a regression of log applications per posting on log job seekers per posting would yield a perfect fit with an elasticity of one. In fact, our cell-level data conform closely to this property (Panel B of Table 7). Moreover, the results in Column (3) of Panel A imply that a two standard deviation increase in slack shrinks posting durations by 2(0.66)(0.030) = 4.0 log points, nearly the same value as before. Thus, our alternative theory-guided slack measure yields the same conclusion: Posting durations exhibit almost no tendency to lengthen as markets tighten.

In general, posting (and vacancy) durations could depend on both job seekers per posting and applications per posting, with separate marginal effects for each. Thus, Columns (5) and (6)

[17] As before, we calculate mean vacancy durations using the method of Davis, Faberman and Haltiwanger (2013), multiplying by (7/6) to convert working days to calendar days.
[18] Fitting column (1) in Table 7 by client type yields an estimated posting-duration elasticity of −0.016 for Recruitment & Staffing firms and −0.095 for Direct Hire firms, as compared to −0.039 in the pooled sample.
[19] Thanks to Pieter Gautier for explaining this feature of their model to us.

report results for specifications that include both slack measures. The elasticity coefficients are again small and, as before, the specification with controls for skill and time fixed effects implies that posting durations actually rise slightly with market slack.[20]

To sum up, our results in Table 7 show that the duration of the posting stage (when job seekers make contact) exhibits almost no tendency to lengthen with market tightness. Many other studies find that vacancy durations rise strongly with market tightness, as we have discussed. The implication is that screening, selection, and other recruitment activities account for the cyclical variation in vacancy durations. Thus, to rationalize our evidence with the implication of MP models that vacancy durations lengthen with market tightness, it is necessary to map the vacancy spell concept in these models to the full set of activities encompassed in the hiring process – including screening, selection, bargaining, and offer acceptance.

8. A Quantitative Sketch of Stages in the Hiring Process

We now draw on several results to create a quantitative sketch of stages in the hiring process. Mean posting duration for job openings on Dice.com is 9.4 days (Table 4). JOLTS data yield a mean vacancy duration of 40.2 calendar days for similar jobs. We combine these pieces of information with evidence from Crane et al. (2016) on the lag between recruitment events and the start of employment by new hires. Their preferred estimate for the mean value of this start lag is 16.2 days.[21] Figure 6 puts this information together and displays it graphically on a timeline that highlights key events and stages in the hiring process. As shown in the figure, the total mean time from date of first posting to the start of employment is 56.4 calendar days.

This depiction captures only the mean duration of each stage in the hiring process. Our results above uncover much heterogeneity in posting durations. Likewise, Crane et al. (2016) find much heterogeneity across recruitment events in the length of start lags. Davis et al. (2013)

---

[20] To obtain the total elasticity of posting durations with respect to slack implied by Column (6), for example, we use the elasticity of applications per posting with respect to job seekers per posting (0.96), and calculate 2[(0.044)(0.45) + (0.96)(-0.017)(0.66)] = 1.8 log points. This quantity says that a two standard deviation increase in job seekers per posting yields a 1.8 log-point *increase* in posting durations when we factor in the associated change in applications per posting.

[21] 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).

document large differences in mean vacancy durations by industry, employer size, employer growth rate, and worker turnover rate. Thus, Figure 6 is best understood as quantifying average outcomes in a process that involves great heterogeneity among employers. Employers may start interviewing and screening applicants on the fly as they continue to gather additional applicants via their posting on Dice.com. Thus, the boundary between the application harvesting stage of the hiring process and the stage devoted to screening, selection and recruitment is a fuzzy one.

IV. Implications for Theories of Search, Matching, and Hiring

We turn now to the implications of our findings for theory and model building, with particular attention to the non-sequential nature of search, the role of labor market intermediaries, the meeting phase of the search and matching process, and the directedness of search. Our discussion leads to further empirical investigations. We also offer some evidence on the extent to which our findings in Dice.com data are indicative of U.S. labor markets more broadly.

1. Employers Do Not Search in the Sequential Manner Posited by Leading Theories

Leading theories of search, matching and hiring posit that employers search sequentially – screening each applicant on arrival, immediately offering a job if the expected match surplus is positive, and ending the search for that position if the offer is accepted. Examples include Diamond (1982), Mortensen (1982, 2003), Mortensen and Pissarides (1994), Pissarides (1985, 2000) and a vast literature that builds on these foundational works. Indeed, modern treatments of frictional unemployment, job-finding rates, job creation incentives, vacancy durations, and wage dispersion in the presence of search and matching frictions are dominated by the sequential search perspective.

There is, however, no general theoretical basis for presuming that sequential search is optimal for employers (or workers). The alternative is a non-sequential strategy, whereby the employer first gathers a pool of applicants, then screens applicants in the pool, selects one or more for an offer, extends the job offer(s), and ends the search for that position if and when the offer is accepted. Employer-side behavior on Dice.com fits this description.[22]

[22] This characterization pertains to standard postings on Dice.com (three-quarters of the total). We cannot say whether such a characterization holds for long-duration postings. Moreover, our evidence pertains to the firm's efforts to fill a given position. Under the linear free-entry condition in many search models, a firm's job creation decision is unaffected by whether it has just filled another position. Our evidence pertains to how employers recruit for a given opening, not how it decides its number of openings.

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. Factors that favor a sequential strategy include a low applicant arrival rate, high costs of screening another applicant, and the absence of scale economies in screening. Factors that favor a non-sequential strategy include a high applicant arrival rate, the bunching of applications shortly after posting, and scale economies in screening.

Figures 2 and 3 and Table 5 show that job seekers target newly posted vacancies. Sixty percent of all applications flow to job openings posted within the first 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 proceeds to screen them. Thus, observed applicant behavior favors non-sequential employer search, according to theory. To put the point starkly, why make an immediate decision about whether to hire the first applicant, if waiting a day or two yields many more applicants? In addition, labor market intermediaries arise partly to exploit scale economies in screening and matching. The prominence of employer-side intermediaries on Dice.com suggests that scale economies are important. That, according to theory, also weighs in favor of non-sequential search strategies on the employer side.

We are not the first to argue that employers often rely on non-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.[23] They also argue that the observed fall in the arrival rate of applicants as postings age, combined with an increase in the job filling rate, is incompatible with sequential search. Their findings align with our evidence and our sketch of the hiring process in Figure 6. Similarly, van Ours and Ridder (1993) find that Dutch employers spend much less time attracting applicants than they devote to evaluating them and selecting one for an offer. This pattern also points to non-sequential search and aligns with Figure 6. Van Ommeren and Russo (2014) argue that sequential employer search implies that the number of rejected applicants is proportional to the number of postings, while non-sequential search implies no such restriction. They reject the proportionality restriction for employers who publicly advertise their vacancies or

[23] Likewise, Barron, Bishop and Dunkelberg (1985) find that "...most employment is the outcome of an employer selecting from a pool of job applicants..." using data from the 1980 Employment Opportunity Pilot Project, a survey of American employers.

rely on employment agencies. Guertzgen and Moczall (2020) report that three-quarters of hires result from a non-sequential search process in a large, representative sample of German employers.

2. Non-Sequential Employer Search Begets Non-Sequential Worker Search

Theoretical reasoning points to a complementarity between non-sequential search on the employer and worker sides of the market. In particular, non-sequential employer search creates a delay between the submission of applications and the selection of a recruit. Given such delays (and hiring delays for other reasons), it makes sense for job seekers to apply for multiple job openings simultaneously while awaiting call-backs and offers, unless applications themselves are very costly to submit. See Morgan and Manning (1985) and Gautier (2002) on this point.[24]

Motivated by this reasoning and our employer-side evidence, we now ask whether job seekers on Dice.com search non-sequentially. To do so, we examine the distributions of applicants and applications by search-spell age at the time of application. A search spell begins when the job seeker applies to a Dice.com posting and has no other applications on the platform in the previous 60 days. Most job seekers have one Dice.com search spell by this definition.

The average number of applications per completed search spell, so defined, is 15.2. In 63 percent of search spells, the jobseeker submits *all* of his or her applications for the entire spell in its first 48 hours. Moreover, half of all applications on Dice.com arise from search spells that are less than two weeks old at the time of application. When restricting attention to first-party applicants (i.e., individuals who don't use third parties to submit applications), two-thirds of them submit all applications in the first 48 hours of their search spell. Their mean number of applications per completed spell is 6.9.[25] These patterns support the theoretical idea that non-sequential employer search begets non-sequential search by job seekers. Faberman and Kudlyak (2019) and Bircini et al. (2025) also find that jobseekers submit multiple applications

3. Why the Distinction between Sequential and Non-Sequential Search Matters

The distinction between sequential and non-sequential search matters for several reasons. First, many workers bargain with prospective employers before accepting a job (Hall and Krueger, 2012). An employer strengthens its bargaining position by gathering a pool of qualified applicants

[24] Burdett and Judd (1983) make essentially the same point to explain why sequential search is not generally superior to non-sequential search in product markets.

[25] When further restricting attention to first-party applications to Direct Hire clients, all such applications occur in the first 48 hours of the search spell for 62 percent of the first-party applicants.

before negotiating with a prospective hire. Likewise, job seekers strengthen their bargaining positions when a non-sequential search strategy yields multiple options. See Appendix E for evidence that the best offer wage rises with the number of offers received by a job seeker. Thus, non-sequential search influences negotiated wage outcomes, which in turn affect search incentives, recruiting behavior, and job creation incentives. Non-sequential search also alters the types of jobs that survive in equilibrium (Galenianos and Kircher, 2009).

Second, in recent theoretical work partly motivated by our evidence, Cai et al. (2025) analyze the determinants of labor market sorting when firms gather a pool of applicants, interview a subset, hire the most profitable interviewee, and then produce. They show that equilibrium allocations depend on both the degree of worker-job production complementarities *and* the number of applicants a firm can interview. Sufficiently strong production complementarities ensure positive assortative matching. Surprisingly though, the degree of complementarity required for positive sorting *rises* in the number of interviews the employer can conduct. In a different model of non-sequential employer search, Birinci et al. (2025) find that lower application costs lead to larger applicant pools, more intensive screening by employers, better matches, and more durable employment relationships. Both papers suggest that non-sequential employer search influences how fully the economy achieves its output potential when sorting and match quality matter.

Third, non-sequential search gives rise to distinct externalities. In the model of Galenianos and Kircher (2009), a worker who accepts a high-wage job offer may also receive and turn down a low-wage job offer, potentially crowding out other applicants for the low-wage job and leaving it unfilled. Job seekers ignore this external effect and submit too many applications from the planner's perspective, and the resulting equilibrium is not constrained efficient.[26] As later work shows, the (in)efficiency of the directed search equilibrium depends on details of the environment.[27] Extra screening costs that applicants impose on employers are another source of inefficiency when job seekers search in a non-sequential manner.

Fourth, Wee (2025) shows that introducing non-sequential search into an otherwise standard search model greatly alters the cyclicality of optimal unemployment insurance (UI) benefits. When jobseekers can submit multiple applications simultaneously, greater worker-side search intensity

[26] Constrained inefficient outcomes also emerge in other equilibrium models with non-sequential search. Examples include Gautier (2002) and Albrecht et al. (2006, 2023).
[27] See Kircher (2009), Gautier and Holzner (2017), Wolthoff (2018), and Wright et al. (2021).

yields more applications per jobseeker, more applicants per vacancy, and greater coordination frictions among firms with respect to which job to offer to which applicant. That, in turn, changes the impact of UI benefit generosity on job creation incentives and the dependence of optimal unemployment benefits on the number of unemployed workers.

In light of these observations, we see our evidence as strong motivation for more attention to models with non-sequential employer search in which one or both sides of the market can simultaneously contact multiple prospective partners before initiating an employment relationship. The quantitative implications of non-sequential search are largely unexplored in the literature.[28]

4. The (Mostly) Overlooked Role of Recruitment and Staffing Firms

Table 6 documents a strikingly large role for Recruitment and Staffing firms on Dice.com. These intermediaries account for two-thirds of postings on the platform and originate more than 60 percent of the applications. We cannot precisely quantify the prevalence of such intermediaries in other U.S. labor markets, but the available evidence suggests their role is large and growing.

Consider staffing firms, which hire workers and lease their labor services to other firms. Staffing firms take on recruiting, screening, and matching functions that would otherwise occur inside their client firms. The staffing-firm share of U.S. payroll employment rose from 1% in 1990 to 2% in 2018.[29] This seemingly modest rise reflects a material shift in how matching happens. Houseman and Heinrich (2015, Table 5) estimate that monthly rates of hires and separations at staffing firms are 7.5 times greater than in the nonfarm economy as a whole. In addition, the employees of staffing firms transition often between client engagements while remaining attached to the staffing firm.[30] Adding client reassignments to hires and separations, the worker reallocation rate of staffing firms is 11.2 times greater than that of other firms. On this basis, Houseman and Heinrich estimate that staffing firms account for 18.5% of all worker reallocation, inclusive of client reassignments, in 2011. Extrapolating from aggregate payroll data, we infer that the staffing-firm share of worker reallocation in the U.S. economy rose from 11% in 1990 to 21% in 2018. This

[28] In a notable exception, Birinci et al. (2025) document that unemployment inflow rates fell by half from 1980 to 2019 in the United States, with little change in the unemployment outflow rate. They show that rising application volumes in a non-sequential search model explain a third of the observed drop in unemployment inflows by leading to better match quality and longer employment relationships.
[29] Calculated as Temporary Help Services employees as a percent of nonfarm payroll employment. The share is 2.1% in the first nine months of 2022.
[30] For their 2007-10 sample period, Houseman and Heinrich find that 58% of worker assignments last less than 1 month, and another 20% last one to three months.

inference accords with independent evidence on the disappearance of short-duration employment relationships and a secular fall in the measured pace of worker reallocation. See, for example, Davis et al. (2010), Hyatt and Spletzer (2013, 2017), Davis and Haltiwanger (2015), and Crump et al. (2019). What once appeared in standard data sources as short-duration employment relationships and between-employer transitions now occurs inside staffing firms.

Several pieces of evidence point to a large and growing role for firms that specialize in headhunting, talent sourcing, screening and other recruitment services for clients that hire employees on their own account. G2, a peer-to-peer business review site, offers ratings and descriptions for more than 150 of the "Best Recruitment Agencies."[31] CareerBuilder.com, Indeed.com, Monster.com and Upwork, among others, have evolved from simple job boards to multi-faceted online platforms that supply talent-sourcing, screening, and recruitment services.[32] Professional networking platforms like LinkedIn and data analytics firms like Lightcast have also evolved to offer sourcing, screening, and recruitment services.[33] In short, recruitment firms and other businesses that provide recruitment services play a major role in U.S. labor markets. They have become more prevalent, partly as an outgrowth of the rise in online job boards.

There are sound reasons to think that staffing and recruitment firms affect matching and other labor market outcomes. Where scale economies in search, recruitment, and screening activities are important, intermediaries can lower the costs of finding prospective workers, assessing their skills, hunting for suitable jobs, and identifying high-quality matches. Their high-volume market engagement gives them better information about job availabilities, suitable workers, potential matches, and likely match quality. Staffing and recruitment firms also have reputational incentives to supply high-quality information and recommendations, in line with the evidence in Stanton and Thomas (2016). Finally, because they can quickly gather a pool of suitable applicants, recruitment firms raise the appeal of non-sequential employer search. We summarize some other evidence of how intermediaries operate on job boards in Appendix E.

---

[31] See https://www.g2.com/categories/recruitment-agencies, accessed on 23 September 2021.
[32] See https://hiring.careerbuilder.com/resume-search, https://hiring.monster.com/products/, https://hiring.monster.com/solutions/recruiting-services/, www.upwork.com/staffing/, and www.indeed.com/hire?co=US&hl=en&from=gnav-menu-homepage, accessed on 24 September. See chapter 2 in Mckinsey Global Institute (2015) on the evolution of online job boards.
[33] See https://business.linkedin.com/talent-solutions/recruiter, https://www.burning-glass.com/solutions/recruiting-and-staffing and https://www.burning-glass.com/products/lens-suite/, accessed on 24 September 2021.

Despite their growing role, theorizing about staffing and recruitment firms is scarce – especially in the form of equilibrium models that speak to frictional unemployment, job-finding rates, job creation incentives, and wage dispersion. In an early effort to model labor market intermediaries, Bull et al. (1987) show that recruitment firms can diversify idiosyncratic risks by sampling over a greater number of job candidates, thereby letting employers fill vacancies more quickly and with greater assurance. Biglaiser (1993) models the role of "middlemen" who specialize in quality assessment and re-sell acquired goods at a premium. Although he considers goods markets, his middlemen perform functions similar to those of staffing firms. Gautier (2002) models how intermediaries reduce duplicative screenings, thereby lowering aggregate screening costs and mitigating congestion externalities. Recruitment and staffing firms perform screening functions akin to those of the intermediaries in Gautier's model. Stanton and Thomas (2016) stress the quality certification role of intermediaries on oDesk.com, adapting a model of Tervio (2009). They develop evidence that these intermediaries improve allocative efficiency and raise the wages of high-quality inexperienced workers.

A quantitatively important part of the matching process operates through institutions that standard search models abstract from. While the behavior we document on the Dice.com platform (short posting spells, application bunching at young postings, and tightness-insensitive posting durations) holds for both Direct Hire employers and Recruitment and Staffing firms, we nonetheless see our evidence on the growing role of intermediaries as strong motivation for greater attention to their impact on match formation, match quality, frictional unemployment, vacancy durations, and the durability of new employment relationships.

5. For Employers, the Meeting Phase Is Short and Unresponsive to Tightness

Vacancy durations lengthen with labor market tightness. The prevailing interpretation of this empirical regularity follows from the leading theory of frictional unemployment, as initially developed by Dale Mortensen and Christopher Pissarides: When labor markets tighten, employers contact job seekers at a slower pace. As the contact rate falls, so too does the vacancy fill rate. Moreover, vacancy spells are coterminous with the meeting phase of the search and matching process in this theory. Even in versions of the theory that include a nontrivial screening process and matching decision, the selection phase happens instantaneously. Thus, these theories highlight the contact rate during the meeting phase as the chief determinant of vacancy durations and their sensitivity to tightness, usually measured as the ratio of vacancies to job seekers.

This prevailing interpretation is at odds with our evidence and with some earlier evidence. Using data for Dutch employers, van Ours and Ridder (1992, 1993) find that vacancy durations mainly reflect the selection phase of the hiring process, not the meeting phase ("application period" in their terminology). Likewise, we find that the meeting phase is much shorter than the selection phase. The sheer brevity of posting durations in our data is hard to square with a literal interpretation of models that treat the meeting phase as coterminous with vacancy spells. Moreover, Table 7 provides evidence that posting durations exhibit little sensitivity to market tightness. For these reasons, we think the contact-rate mechanism emphasized by existing theory lacks a persuasive explanation for the duration of vacancy spells *and* for the sensitivity of vacancy durations to labor market conditions.

6. The Non-Random Character of Worker-Side Search

Application flows can depart from a random allocation to postings for multiple reasons. First, job seekers sort themselves into markets defined by skill requirements, job function, location, industry, and more. Even if job seekers search randomly within markets, sorting across markets leads to non-random allocations at more aggregated levels. We will assess how fully this type of sorting rationalizes the non-random allocation of applications to postings.

Second, employers can advertise wages and other job attributes in their postings to further "direct" the flow of applicants. Other things equal, employers that offer higher pay can expect to attract more (and possibly better) applicants. Early theoretical developments of this idea include Montgomery (1991), Shimer (1996, chapter 1), Moen (1997), Mortensen and Wright (2002) and Shi (2002). Wright et al. (2021) review the broader literature. Unlike in traditional search models where wages emerge from a bargaining process, wages in directed search models reflect a competitive process. This form of directed search often yields more efficient allocations than random search, because offer wages help guide the flow of applicants to suitable jobs, and because the commitment aspect of wage posting alleviates hold-up problems.[34] Partly for these reasons, it is important to assess the role of wage posting in directing the flow of applications to job vacancies.

We now fit a series of statistical models to quantify distance from randomness and to assess how job, employer, and posting characteristics affect application flows. Given the high dispersion of applications over postings, we consider negative binomial (NB) models with mean μ and

[34] For early analyses that study hold-up problems in search models with wage bargaining, see Acemoglu (1996) and Davis (2001).

dispersion parameter θ.[35] These parameters pin down the variance of applications per posting at $\mu(1+\theta\mu)$. When $\theta = 0$, the NB model collapses to a Poisson distribution, which approximates the binomial distribution as the posting count gets large while μ remains constant. For $\theta > 0$, the variance exceeds the mean. Thus, the variance-mean ratio, $\frac{\sigma^2}{M} = 1 + \theta\mu$, is a natural metric for distance from randomness in the NB model, as is the value of θ and the gap between the flow-weighted and unweighted means – i.e., $M^W - (M+1)$.

As a starting point, Table 8 reports statistics on the distribution of applications over completed posting spells in the raw data and for two simple statistical models with no covariates. Comparing columns (1) and (2) reinforces the earlier point that random assignment is a poor characterization of the data. Column (3) fits an NB model that targets the mean and standard deviation of applications per posting. The model requires a high value of the dispersion parameter ($\theta = 6.86$) to match the standard deviation of applications per posting in the data. While this simple NB model fits much better than the binomial model, it overstates the observed fraction of postings with no applications by a factor of 2.6.

To generalize the NB model and allow for covariates, let the number of applications to posting *i,* A_i, be a random variable that obeys an NB distribution with mean μ_i and dispersion θ. Let the expected number of applications be $\ln(\mu_i) = \beta_0 + \mathrm{X}_\mathrm{i}'\beta$, where X_i is a vector of job, posting and employer characteristics. The probability that posting i receives exactly a applications is

$$P(A_i = a|X_i) = \frac{\Gamma\left(a+\frac{1}{\theta}\right)}{\Gamma(a+1)\Gamma\left(\frac{1}{\theta}\right)} \frac{(\theta\mu_i)^a}{(1+\theta\mu_i)^{a+\frac{1}{\theta}}} \quad \text{for } a = 0, 1, 2, \ldots \qquad (2)$$

with log-likelihood given by

$$\ell_i(\beta_0, \beta, \theta; a, X) = a_i \log(\theta\mu_i) - \left(a_i + \frac{1}{\theta}\right)\log(1+\theta\mu_i) + \log\left(\frac{\Gamma\left(a_i+\frac{1}{\theta}\right)}{\Gamma\left(\frac{1}{\theta}\right)}\right), \qquad (3)$$

where $\Gamma(\cdot)$ is the Gamma function.

Figure 7 summarizes the performance of (2) and (3) when fit by maximum likelihood.[36] Point 1 in the upper right describes the raw data, with a flow-weighted mean number of applications

[35] Here, we drop postings with no views. In an earlier draft, we retained these postings and fit zero-inflated negative binomial models to accommodate them. That model is considerably more complex but yields very similar results. For simplicity, we stick to NB models in our reported results.

[36] We use the GENMOD procedure to estimate the model, as set forth in SAS Institute Inc. (2016) and as based on Cameron and Trivedi (1998).

per posting of 83.6 and a standard deviation of applications per posting of 27.7. This point corresponds to the model (2) and (3) with $\log(\mu_i) = \beta_0$. Point 6 in the lower left depicts the random assignment benchmark, with a flow-weighted mean of 11.5 and standard deviation of 3.2. The other points tell us how well various explanatory variables account for the distribution of application flows to postings and departures from the random assignment benchmark.

Consider Point 2, which reflects an X vector with fixed effects for 54 skill categories and 54 job locations (Metropolitan Statistical Areas, or MSAs) and a continuous control for completed spell duration. The X vector also includes 71 monthly time effects to control for fluctuations in overall market tightness and other sources of systematic time variation in applications per posting. This model yields an estimated θ of 1.59, much closer to the random assignment benchmark. We conclude that the self-sorting of job seekers into specific labor markets defined by skill and MSA is a major source of unevenness in the allocation of applications to postings.

That said, this statistical model is far from the random assignment benchmark. To see this point, note that the model-implied value of the flow-weighted mean within labor markets defined by skill and MSA is about 30, nearly three times the simple mean of 10.5 implied by random assignment. The model-implied standard deviation within markets is four times the value implied by random assignment. Thus, the self-sorting of job seekers into markets defined by skill and job location leaves much of the unevenness in applications per posting unexplained.

Expanding the X vector to include posting characteristics yields Point 3 in Figure 7 and the estimates reported in Column (1) of Table 9. According to these results, postings that accept third-party applications attract about 225 percent more applicants on average, conditional on the other variables in the statistical model.[37] This large effect aligns with our earlier finding that third parties submit most of the applications on Dice.com. The results in Column (1) also say that positions posted by Staffing & Recruitment firms attract 23 percent fewer applicants, on average, and postings that accommodate submissions directly on Dice.com attract 16 percent more. The estimated θ is 1.28. All of these parameters are precisely estimated. Thus, posting characteristics help explain the unevenness of application flows, but this version of the generalized NB model is also far from random assignment.

[37] Computed as $100[\exp(1.1.8) - 1]$, using the results in column (1) of Table 9.

Next, we turn to the role of posted wages in directing application flows. To do so, we further expand the X vector to include indicators of whether pay is specified in hourly, weekly, or annual terms. We also interact each pay-period indicator with a continuous measure of log pay per unit time. Here, the omitted category covers postings that do not specify an offer wage or wage range. Remarkably, 83 percent of postings do not state a pay level or range. This simple result undercuts the view that posted wages play a central role in directing applicants to job vacancies or in preventing hold-up problems that arise with bargaining after search. Batra et al. (2023) report that 86 percent of online vacancy postings in the United States from 2012 to 2017, as captured by Lightcast, do not state a pay level or range. So, the result extends beyond the Dice.com platform.

Perhaps, however, stated offer wages in some postings play a material role in directing application flows on Dice.com. Point 4 in Figure 7 and column (2) in Table 9 speak to this matter. As it turns out, the wage-posting measures do not influence the direction of application flows. To see this point, note that adding these variables to the X vector has no discernable impact on the estimated value of θ, nor does it alter the implied conditional values for the flow-weighted mean and the standard deviation of applications per posting within markets.[38] Moreover, the coefficient estimates on the wage variables in column (2) are miniscule in magnitude and precisely estimated. To explore the robustness of this result, we refit the models by skill category, letting all model parameters vary freely with skill requirements. Here as well, we find that the wage variables play almost no role in directing the flow of applicants to particular vacancies or in explaining departures from random assignment. Table C.3 reports results separately for the top twelve skill categories on Dice.com, as measured by number of postings. In short, explicitly posted wages play almost no role in directing the flow of applicants to particular vacancies.

The rest of Table 9 considers even more expansive X vectors that include about 1,600 job title fixed effects in Column (3) and about 4,100 employer-specific fixed effects in Column (4). Conceptually, these models allow for an even more granular self-sorting of job seekers into narrowly defined labor markets. As seen in Table 9 and in Point 5 of Figure 7, these additional variables help explain the direction of application flows. Nevertheless, the model remains some

[38] When we restrict the sample to postings with a stated offer wage or wage interval, the estimated value of θ is 1.269 for the Column (1) specification and 1.268 for Column (2). In addition, the coefficients on the ln(wage) variables are tiny, ranging from -0.01 to 0.02.

distance from the random assignment benchmark. The same is true when we fit models with highly expansive X vectors separately by skill category, as shown in in Table C.3.

Summing up, job seekers sort themselves across labor markets in a highly uneven manner. We can largely, though not entirely, account for this unevenness using statistical models that relate the direction of application flows to observed job, posting, and employer characteristics. To our surprise, however, posted wages exert no influence on the distribution of application flows to postings on the Dice.com platform. This result undercuts the central premise in search theories founded on the idea that explicitly posted wages direct job seekers. It leaves open the possibility that application flows are directed by expected compensation, job amenities, or other attributes of the posting that job seekers infer from the job title or the employer's identity. Finally, even our most expansive and flexible statistical models imply large departures from a random allocation benchmark. Perhaps other – as yet unobserved – attributes of jobs, postings, employers, and workers can fully explain the observed departures from random assignment.

7. Stock-Flow Matching

Our evidence that job seekers favor newly posted vacancies points to the empirical relevance of stock-flow matching theories. See Coles and Smith (1998), Gregg and Petrongolo (2005) and Ebrahimy and Shimer (2010) for prominent examples. As we have stressed, however, our results also underscore the importance of non-sequential employer search. And, as highlighted by Figure 6, time devoted to screening, interviewing, selection and negotiation – activities that come after the meeting phase of the matching process – largely account for the time it takes to fill an open job position. Waiting for a new inflow of job seekers before matching can commence, as envisioned in some stock-flow matching models, appears to play a minor role on the Dice.com platform.[39]

V. Concluding Remarks

This paper examines application flows and vacancy postings on Dice.com, a platform for jobs and workers in software design, computer systems, engineering, financial analysis, management consulting, and other occupations that require technical skills. Some of our findings

[39] This pattern could reflect greater jobseeker "exposure" to younger postings. Appendix Figure C.4 shows that applications per view fall with posting age. That is, a job seeker who views an older posting is less likely to apply to it than one who views a new posting. Thus, the concentration of application flows at young postings at least partly reflects job-seeker behavior conditional on exposure. The Dice.com platform did not preferentially display younger postings during our sample period.

challenge leading search theories or how those theories are commonly interpreted. Other findings highlight the understudied role of intermediaries in the search and matching process.

One challenge pertains to our evidence on the prevalence of non-sequential search behavior. In contrast, leading theories of frictional unemployment and vacancy durations presume that employers and workers search in a sequential manner. Economic reasoning and theoretical work imply that sequential and non-sequential search differ in their implications for wages, search incentives, job creation incentives, screening intensity, match quality, the durability of employment relationships, and the optimal design of unemployment insurance. In light of these observations, we see our evidence as strong motivation for greater attention to theories that feature non-sequential search. Our evidence also calls for more empirical research into the choice of search strategies, and how and why that choice varies across labor markets, institutional settings, and time.

Our study also underscores the limitations of theories that focus exclusively on the meeting phase of the search and matching process. Specifically, we find that screening, interviewing, selection and negotiation activities largely account for the duration of vacancies and the cyclicality of vacancy fill rates. The prevailing theoretical explanations for vacancy durations and cyclicality in fill rates rest on claims about how market tightness affects the meeting rate. Those explanations are untenable for the vacancy postings on Dice.com, because the meeting phase is so brief and because its duration varies so little with tightness. Thus, we see our evidence as calling for a re-examination of what determines vacancy durations and why fill rates fluctuate over time.

Another challenge relates to the highly uneven distribution of applications over vacancy postings. We can explain most, but not all, of this unevenness as the outcome of applicant self-sorting across labor markets defined by job location, skill requirements, job characteristics, and employer identity. Surprisingly, however, offer wages in the vacancy postings play essentially no role in directing application flows or rationalizing departures from a random allocation. This finding undercuts a central premise in search theories that treat posted wages as a key influence on the direction of application flows.

To be sure, other information in vacancy postings can influence the direction of application flows. Nothing in our study refutes the idea that job seekers form expectations about wages from the job location, skill requirements, job title and other information stated in vacancy postings. However, that type of information does not head off the transaction costs associated with bargaining over pay after contact, screening, interviews, selection, etc. Nor does it offer the kind of commitment that

alleviates hold-up problems that can arise when parties to a prospective match make specific match-relevant investments before entering into an employment relationship.

We uncover a huge role for other intermediaries that operate on the Dice.com platform. Recruitment and staffing firms account for two-thirds of all postings and attract more than sixty percent of the applications. In addition, Dice.com itself provides a range of intermediary services by regulating the applicant pool, letting employer-side clients screen out third-party applicants, giving them access to high-quality résumé banks, letting them ping workers to alert them to specific postings, and by improving worker-side search functionality over time. Perhaps intermediaries play a large role on Dice.com because the jobs posted there involve well-defined technical skills. Addressing that issue requires the study of other platforms.

We also provide evidence that recruitment and staffing firms, and intermediaries more broadly, have come to play an increasingly important role in U.S. labor markets in recent decades. Theories of labor market intermediaries are relatively scarce. Theoretically grounded quantitative analyses of search and matching with a prominent role for labor market intermediaries are scarcer yet.

References

Acemoglu, Daron, 1996. "A Microfoundation for Social Increasing Returns to Human Capital Accumulation," *Quarterly Journal of Economics,* 111 (August), 779-804.

Albrecht, James, Xiaoming Cai, Pieter Gautier and Susan Vroman, 2020. "Multiple Applications, Competing Mechanisms, and Market Power," *Journal of Economic Theory,* 190 (November).

Albrecht, James, Bruno Decreuse and Susan Vroman, 2023. "Directed Search with Phantom Vacancies," International Economic Review, 64, no 2 (May), 837-869.

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

Banfi, Stefano and Benjamin Villena-Roldán, 2019. "Do High-Wage Jobs Attract More Applicants? Directed Search Evidence from the Online Labor Market," *Journal of Labor Economics*, 37, no. 3, 715-746.

Barron, John M., John Bishop and William C. Dunkelberg, 1985. "Employer Search: The Interviewing and Hiring of New Employees," *Review of Economics and Statistics,* 67, no. 1.

Batra, Honey, Amanda Michaud and Simon Mongey, 2023. "Online Job Posts Contain Very Little Wage Information," NBER Working Paper 31984.

Bellemare, Marc F. and Casey J. Wichman, 2020. "Elasticities and the Inverse Hyperbolic Sine Transformation," *Oxford Bulletin of Economics and Statistics* 82, no. 1 (February), 50-61.

Birinci, Serdar, Kurt See and Shu Lin Wee, 2025. "Job Applications and Labor Market Flows," *Review of Economic Studies,* 92, no. 3 (May), 1438-1496.

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 Kennth L. Judd, 1983. "Equilibrium Price Dispersion," *Econometrica,* 51, no. 4 (July), 955-969.

Cai, Xiaoming, Pieter Gautier and Ronald Wolthoff, 2025. "Search, Screening and Sorting," *American Economic Journal: Macroeconomics*, 17, no. 3 (July), 205-236.

Cameron, A. Colin and Pravin K. Trivedi, 1998. *Regression Analysis of Count Data.* Cambridge: Cambridge University Press.

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

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

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, University of Chicago.

Crump, Richard K., Stefano Eusepi, Marc Giannoni and Aysegul Sahin, 2019. "A Unified Approach to Measuring *u**," *Brookings Papers on Economic Activity,* Spring.

Davis, Steven J., 2001. "The Quality Distribution of Jobs and the Structure of Wages in Search Equilibrium," NBER Working Paper No. 8434.

Davis, Steven J., R. Jason Faberman and John C. Haltiwanger, 2010. "Business Volatility, Job Destruction, and Unemployment," *American Economic Journal: Macroeconomics,* 2 (April).

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., R. Jason Faberman, John C. Haltiwanger and Ian Rucker, 2009. "Adjusted Estimates of Worker Flows and Job Openings in JOLTS." In *Labor in the New Economy*, edited by Katharine Abraham, Michael Harper and James Spletzer.

Davis, Steven J. and John Haltiwanger, 2015. "Labor Market Fluidity and Economic Performance," NBER Working Paper No. 20479. Federal Reserve Bank of Kansas City Jackson Hole Symposium on *Re-Evaluating Labor Market Dynamics.*

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, 2019. "The DHI Vacancy and Application Flow Database: Record Layouts, Variable Descriptions, and Basic Summary Statistics," Technical Memorandum, 2 February.

Davis, Steven J. and Brenda Samaniego de la Parra, 2021. "Platform Design and Application Flows: A Case Study," work in progress.

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

Diamond, Peter A., 1982. "Wage Determination and Efficiency in Search Equilibrium," *Review of Economic Studies,* 49, no. 2 (April), 217-227.

Ebrahimy, Ehsan and Robert Shimer, 2010. "Stock-Flow Matching," *Journal of Economic Theory,* 145, no. 4 (July), 1325-1353.

Faberman, R. Jason and Marianna Kudlyak, 2019. "The Intensity of Job Search and Search Duration," *American Economic Journal: Macroeconomics,* 11, no. 3 (July), 327-357.

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, 2017. "Simultaneous Search and Efficiency of Entry and Search Intensity," *American Economic Journal: Microeconomics,* 9, no. 3, 245-282.

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.

Guertzgen, Nicole and Andreas Moczall, 2020. "Sequential versus non-sequential search among German employers – evidence from a job vacancy survey," *Applied Economics Letters*, 27:11, 873-879.

Hall, Robert E. and Alan B. Krueger, 2012. "Evidence on the Incidence of Wage Posting, Wage Bargaining, and the On-the-Job-Search," *American Economic Journal: Macroeconomics,* 4, no. 4 (October), 56-67.

Houseman, Susan N. and Carolyn J. Heinrich, 2015. "Temporary Help Employment in Recession and Recovery," Upjohn Institute Working Paper 15-227.

Hyatt, Henry R. and James R. Spletzer, 2013. "The Recent Decline in Employment Dynamics," *IZA Journal of Labour Economics,* 2, no. 5.

Hyatt, Henry R. and James R. Spletzer, 2017. "The Recent Decline of Single Quarter Jobs," *Labour Economics,* 46 (June), 166-176.

Kircher, Philipp, 2009. "Efficiency of Simultaneous Search," *Journal of Political Economy,* 117, no. 5 (October), 861-913.

Leduc, Sylvain and Zheng Liu, 2020. "The Weak Job Recovery in a Macro Model of Search and Recruiting Intensity," *American Economic Journal: Macroeconomics,* 12, no. 1, 310-343.

Marinescu, Ioana and Ronald Rathelot, 2018. "Mismatch Unemployment and the Geography of Job Search," *American Economic Journal. Macroeconomics*, 10, no. 3, 42-70.

Marinescu, Ioana and Ronald Wolthoff, 2020. "Opening the Black Box of the Matching Function: The Power of Words" *Journal of Labor Economics*, 38, no. 2, 535-568.

McKinsey Global Institute, 2015. *A Labor Market that Works: Connecting Talent with Opportunity in the Digital Age.* McKinsey & Company.

Modestino, Alicia Sasser, Daniel Shoag, and Joshua Balance, 2020. "Upskilling: Do Employers Demand Greater Skill When Workers Are Plentiful?" *Review of Economics and Statistics*, 102, no. 4, 793-805.

Moen, Espen R., 1997. "Competitive Search Equilibrium," *Journal of Political Economy,* 105, no. 2, 385-411.

Mongey, Simon and Gianluca Violante, 2025. "Macro Recruiting Intensity from Micro Data," forthcoming in the *American Economic Journal: Macroeconomics*.

Montgomery, James D., 1991. "Equilibrium Wage Dispersion and Interindustry Wage Differentials," *Quarterly Journal of Economics,* 106, no. 1 (February), 163-179.

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.

Mortensen, Dale T. and Randall Wright, 2002. "Competitive Pricing and Efficiency in Search Equilibrium," *International Economic Review,* 43, 1-20.

Mueller, Andreas I, Damian Osterwalder, Josef Zweimüller and Andreas Kettemann, 2024. "Vacancy Durations and Entry Wages: Evidence from Linked Vacancy-Employer-Employee Data," *Review of Economic Studies*, 91, no. 3, 1807-1841.

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.

Rabinovich, Stanislav, Brenda Samaniego de la Parra and Ronald Wolthoff, 2024, "Wage Setting Protocols and Labor Market Conditions: Theory and Evidence"

SAS Institute Inc., 2016. *SAS/STAT® 14.2 User's Guide. Generalized Linear Models Theory*. Cary, NC: SAS Institute Inc.

Shi, Shouyong, 2002. "A Directed Search Model of Inequality with Heterogeneous Skills and Skill-Biased Technology," *Review of Economic Studies,* 69, no. 2, 467-491.

Shimer, Robert, 1996. *Essays in Search Theory.* Massachusetts Institute of Technology, Department of Economics.

Stanton, Christopher T. and Catherine Thomas, 2016. "Landing the First Job: The Value of Intermediaries in Online Hiring," *Review of Economic Studies,* 83, no. 2 (April), 810-854.

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

van Ommeren, Jos and Giovanni Russo, 2014. "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.

Wee, Shu Lin, 2025. "Optimal Unemployment Insurance with Multiple Applications," working paper, University of Melbourne.

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

Wright, Randall, Philipp Kircher, Benoit Julien and Veronica Guerrieri, 2021. "Directed Search and Competitive Search Equilibrium: A Guided Tour," *Journal of Economic Literature,* 59, no. 1. (March), 90-148.

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

	Total, Millions	Direct Hire, Millions	Recruitment and Staffing Firms, Millions
(1) Number of Raw Vacancy Postings	7.5	2.5	5.0
(1.a) Standard Postings	5.6	1.7	3.9
(1.b) Long-Duration Postings	1.9	0.8	1.1
(2) Volume of Applications	125.3	47.9	77.4
(2.a) Email Applications	95.3	34.4	60.9
(2.b) URL Applications	30.0	13.4	16.6

Notes: "Direct Hire" clients 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 distinct Job ID values in the DHI Database, and Row (2) reports the number and distribution of applications. "Standard Postings" are those for which the client permanently removes the posting with 30 days (720 hours) after first posting. "Email Applications" refer to ones submitted directly on the Dice.com platform, and "URL Applications" refer to the frequency with which job seekers click through to an external URL.

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.

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 Applications
250 Job IDs	1,285	93.5%	95.2%
100 Job IDs	1,983	95.0%	96.5%
50 Job IDs	2,746	95.7%	97.1%

Notes: Column (2) reports the number of distinct job titles that meet the minimum posting frequency specified in Colum (1). Columns (3) and (4) report the shares of Job IDs and Applications accounted for by these frequently posted job titles.

Table 4. The Distribution of Completed Posting Durations by Job Function and Application 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
By Selected Job Functions							
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 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. 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 vacancy postings with first posting date on or after December 1, 2017 to avoid the inclusion of incomplete spells.

Table 5. Selected Statistics on Applications Per Posting

	Mean Applications Per Vacancy		*Percent of Applications Received Within:*	
	Equal Weighted	Flow Weighted	First 48 Hours After Posting	First 96 Hours After Posting
	(1)	(2)	(3)	(4)
All Standard Postings	11.0	88.1	45.3	59.9
All Job Titles with at Least 100 Standard Postings	11.2	89.1	45.3	60.0
Job Titles Sorted by Mean Applications Per Posting				
Bottom Quintile	3.4	19.8	32.9	48.3
Fourth Quintile	5.4	30.7	38.1	52.9
Third Quintile	7.4	41.9	38.6	53.7
Second Quintile	10.8	60.4	44.8	59.7
Top Quintile	22.3	134.6	49.8	64.0
Selected Job Functions				
Developer	16.3	141.3	49.9	64.2
Engineer	7.5	64.4	40.9	55.7
Administrator	10.4	58.5	45.3	60.2
Mechanical Engineer	4.1	17.5	26.0	41.1
Electrical Engineer	3.7	15.2	24.4	40.6
Business Analyst	22.5	97.0	49.5	63.1
Analyst	9.9	67.4	39.5	54.4
Help / Support Desk	7.5	32.5	29.7	45.4
Sales / Business Development	3.0	24.0	28.5	43.6

Notes: Except for the first row, entries pertain to standard postings with completed spells and at least 100 postings. Columns (3) and (4) report flow-weighted statistics for postings that receive at least one application. Equal-weighted statistics are quite similar (within a given row). In constructing this table, we dropped vacancy postings with first posting date on or after December 1, 2017 to avoid the inclusion of incomplete spells.

Table 6. Intermediaries Dominate Activity on Both Sides of the Dice.com Platform

The Joint Distribution of Applications over Employer-Side Client Types and Worker-Side Application Types, 2015 to 2017			
	1st-Party Applications	3rd-Party Applications	Not Classified
Direct Hire Clients	12%	22%	3%
Recruitment & Staffing Firms	20%	39%	4%

Notes: This table restricts attention to applications from 2015 to 2017, because the DHI Database does not distinguish between 1st-party and 3rd-party applications before 2015.

Table 7 How Posting Durations Vary with Market Slack, Monthly Data and 48 Skill Categories

Panel A. Main Regressions

Dependent Variable: ln(Mean Duration of Postings in Skill Category *j* in Month *t*)

	(1)	(2)	(3)	(4)	(5)	(6)
ln(Job Seekers/Postings)	-0.039***	0.027***			0.013	0.044***
	(0.003)	(0.005)			(0.008)	(0.008)
ln(Applications/Postings)			-0.030***	0.009**	-0.039***	-0.017***
			(0.002)	(0.004)	(0.006)	(0.006)
Constant	2.34***	2.28***	2.37***	2.28***	2.22***	2.15***
	(0.003)	(0.011)	(0.005)	(0.012)	(0.012)	(0.015)
Observations	3,408	3,408	3,408	3,408	3,408	3,408
R-squared	0.64	0.74	0.64	0.74	0.64	0.74
Within R-squared	0.05	0.01	0.06	0.001	0.06	0.01
Fixed Effects	Skill	Skill & Time	Skill	Skill & Time	Skill	Skill & Time

Panel B. Auxiliary Regressions

Dependent Variable: ln(Applications Per Posting in Skill Category *j* in Month *t*)

	Elasticity of Applications Per Posting with Respect to Job Seekers Per Posting	Within R-squared
Controlling for Skill Fixed Effects	1.36 (0.009)	0.87
Controlling for Skill and Time Fixed Effects	0.96 (0.014)	0.58

Panel C. Selected Summary Statistics

	Log Posting Durations	Log Job Seekers Per Posting	Log Applications Per Posting
Standard Deviation Across Skill-Month Cells	0.14	0.61	0.83
Average Standard Deviation over Time within Skill Categories	0.08	0.45	0.66

Notes: The sample contains monthly data from January 2012 to November 2017 for all 48 skill categories that attract at least 25 postings in every month. We group postings into skill-by-month cells based on the first skill requirement mentioned in the extended job title and the month in which the posting first became active. Within each cell, we calculate the mean posting duration as the ratio of total posting days to the number of postings, where total posting days is the time elapsed from first to last active date summed over all active postings in the cell. Job Seekers per Posting equals the number of distinct applicants with an application to any posting in the cell divided by the number of postings in the cell. Applications per Posting equals the number of distinct applications to any posting in the cell divided by the number of postings in the cell. The Within R-squared is for the regression that first sweeps out the indicated fixed effects. To obtain the elasticity of applications per posting with respect to job seekers per posting, we regress the former on the latter in the cell-level data while controlling for the indicated fixed effects.

Table 8. Selected Statistics for Standard Postings with Completed Spells

	(1)	(2)	(3)
	DHI Data	Binomial Model (Random Assignment): μ=10.5	Negative Binomial Model: μ=10.5, θ=6.86
Simple Mean of Applications per Posting	10.5	10.5	10.5
Standard Deviation of Applications per Posting	27.7	3.2	27.7
Percent of Postings with No Applications	20.4	0.003	53.5
Flow-Weighted Mean of Applications per Posting	83.6	11.5	83.6
Ratio of Flow-Weighted to Simple Mean	8.0	1.1	8.0
Ratio of Variance-to-Mean	73.1	1.0	73.1

Notes: Column (1) reports statistics for the raw data, and columns (2) and (3) report them for the indicated models. We target the simple mean of applications per posting when fitting the Binomial Model, and the simple mean and standard deviation when fitting the negative binomial Model. The sample contains 48,812,063 applications to 4,785,218 job vacancies first posted on Dice.com from 1 January 2012 to 30 November 2017. We exclude vacancies first posted on or after 1 December 2017 to focus on completed spells. We exclude postings with zero views and postings that mention none of our skill categories.

Table 9. Negative Binomial Models with Covariates

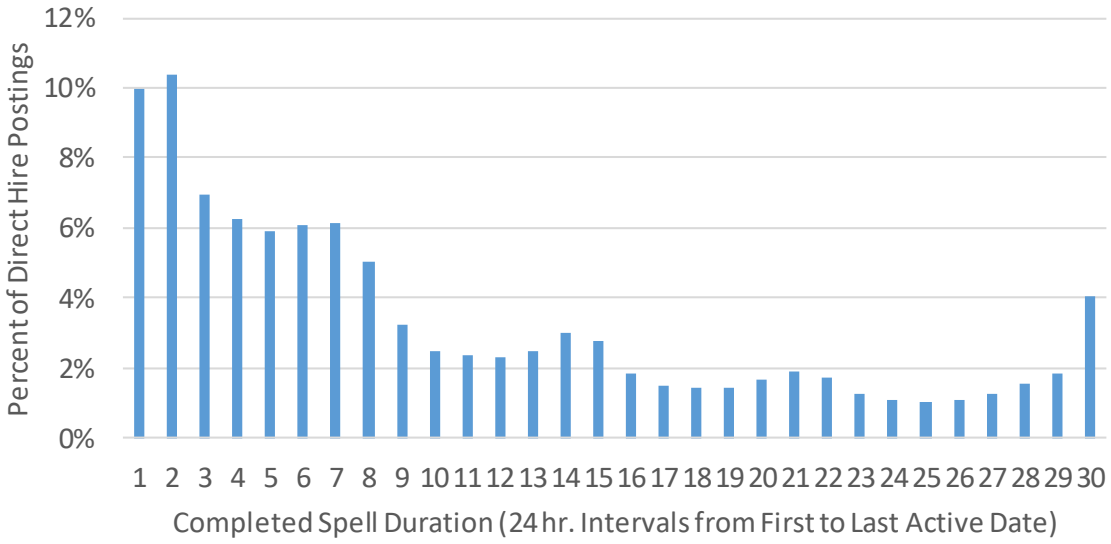
Unit of Analysis: Completed Posting Spell
Dependent Variable: Number of Applications to the Posting

Type and Number of Covariates →		(1) Fixed effects for 54 Skill Categories, 71 Time Periods, and 54 MSAs	(2) + 3 pay-period indicators and 1 pay variable for each pay period	(3) +1,626 Job Title Fixed Effects	(4) +4,115 Employer-Specific Fixed Effects
Pay period	Hourly		0.02 (0.007)	-0.03 (0.006)	0.01 (0.007)
	Weekly and Other		-0.14 (0.028)	-0.02 (0.027)	0.00 (0.027)
	Annual		-0.01 (0.008)	-0.02 (0.007)	-0.01 (0.008)
ln(wage) by pay-period category	Hourly		0.00 (0.002)	0.01 (0.004)	0.00 (0.004)
	Weekly and Other		0.03 (0.006)	0.00 (0.002)	0.00 (0.003)
	Annual		-0.01 (0.002)	-0.01 (0.002)	0.00 (0.002)
3rd-Party Applications Are Okay		1.18 (0.002)	1.18 (0.002)	1.20 (0.002)	1.18 (0.002)
Posted by Recruitment & Staffing Firm		-0.26 (0.002)	-0.26 (0.002)	-0.29 (0.002)	-0.40 (0.007)
Applications Submitted on Dice.com		0.15 (0.002)	0.15 (0.002)	0.16 (0.002)	0.29 (0.003)
Dispersion Parameter, θ		1.28 (0.001)	1.27 (0.001)	1.14 (0.001)	0.98 (0.001)

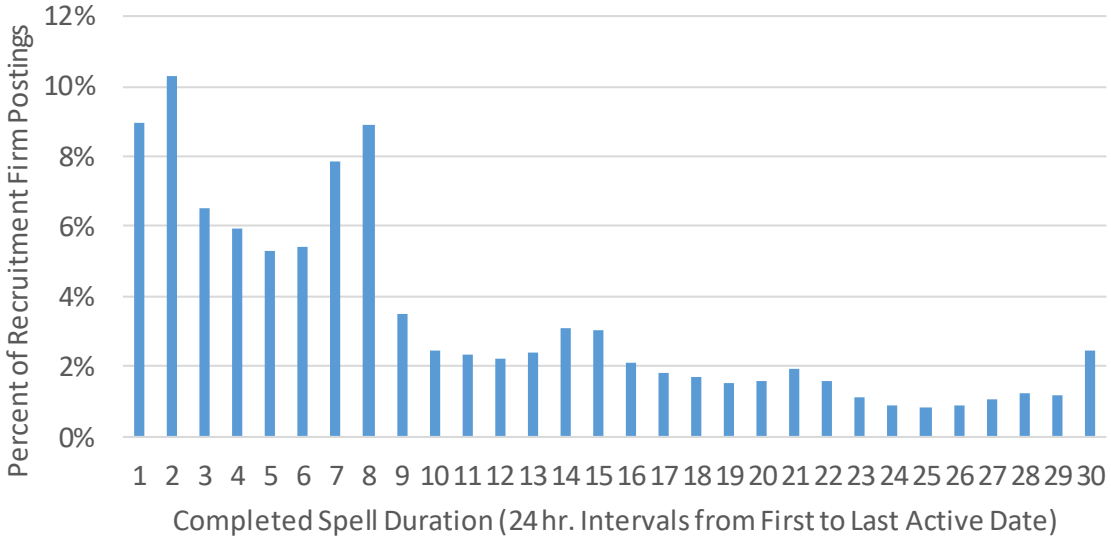
Notes: Columns (1) to (4) report results for versions of model (2) and (3) with increasingly expansive X vectors, as indicated in the column headings. The sample is the same as in Table 8. We estimate the models by maximum likelihood and report standard errors in parentheses. The X vector includes an indicator variable for each pay-period category and one ln(wage) measure for each category. The baseline "pay period" pertains to the 83% of postings that do not state a pay value or interval. Among postings that state pay, 55% offer an hourly wage and 42% offer an annual salary. The other 3% state pay for a daily, weekly, bi-weekly or monthly period, which we convert to a weekly wage assuming a five-day workweek or 4.5 weeks per month. Among the 17% of posting that contain numeric information about pay, 57% state a pay interval. In these cases, we use the midpoint. The results are robust to instead using the lower end of the pay interval. When the acceptability of third-party applications is not reported (4% of observations), we set the value to Okay. If a posting received no applications, we randomly assign it to an application channel.

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: We remove jobs first posted on or after December 1, to exclude incomplete spells.

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

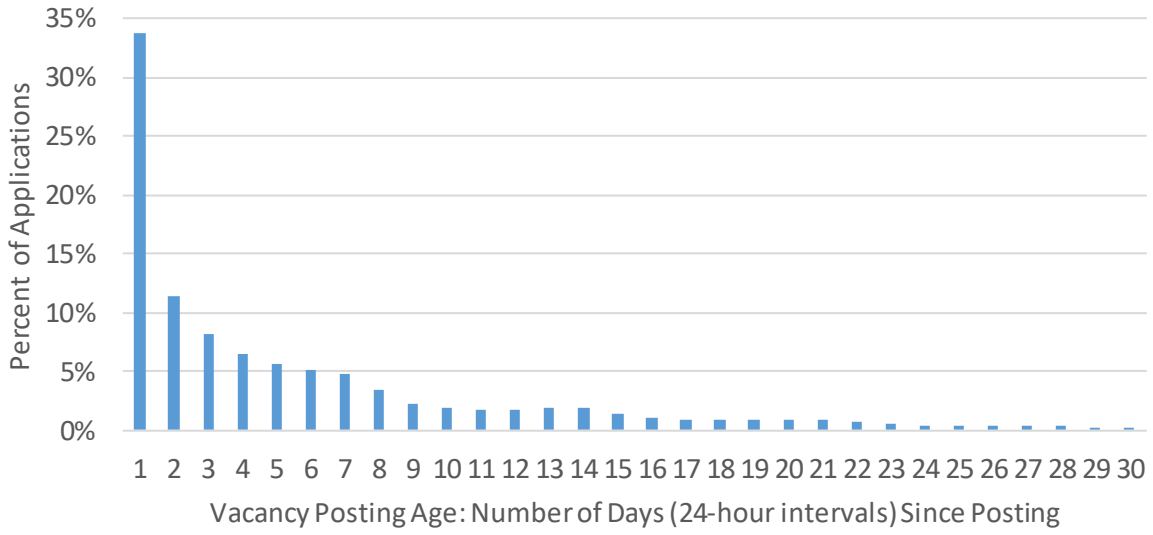
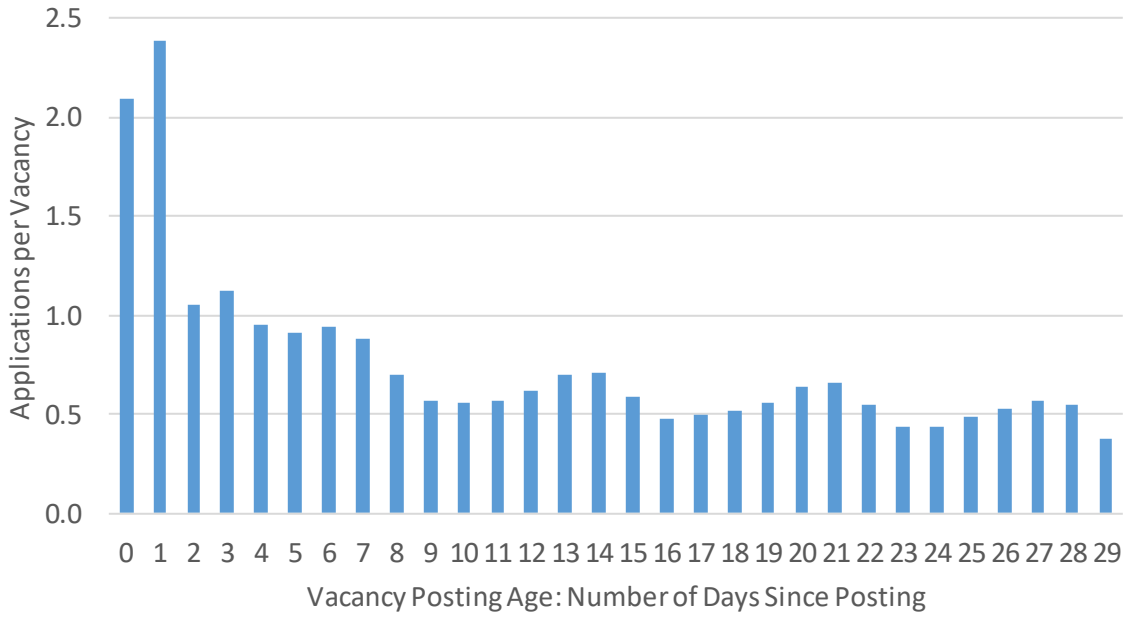


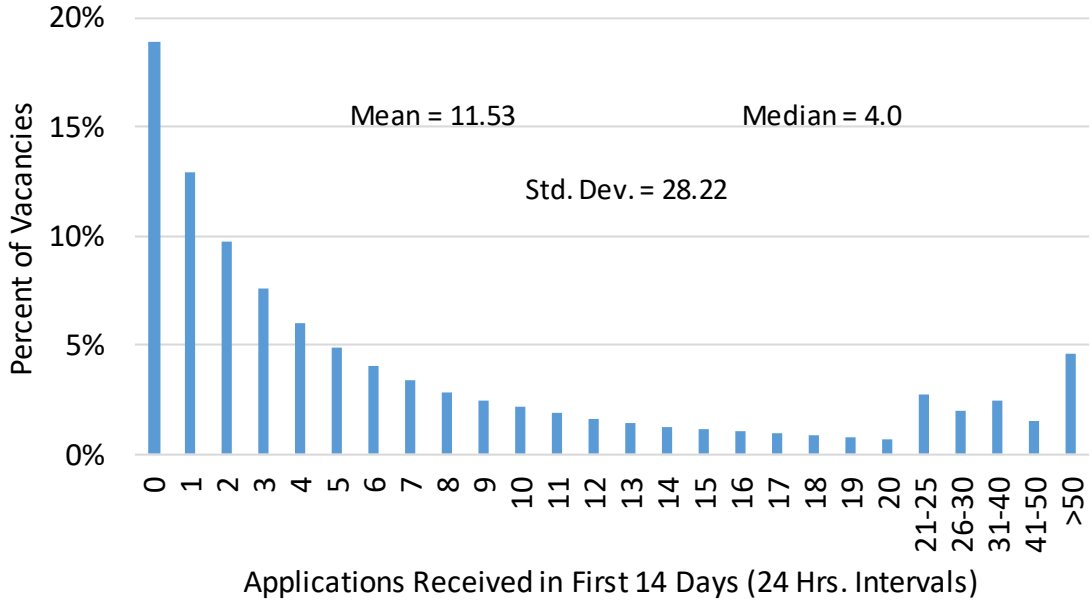
Figure 3. Mean Daily Applications Per Vacancy by Posting Age, Standard Postings, January 2012 to December 2017



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

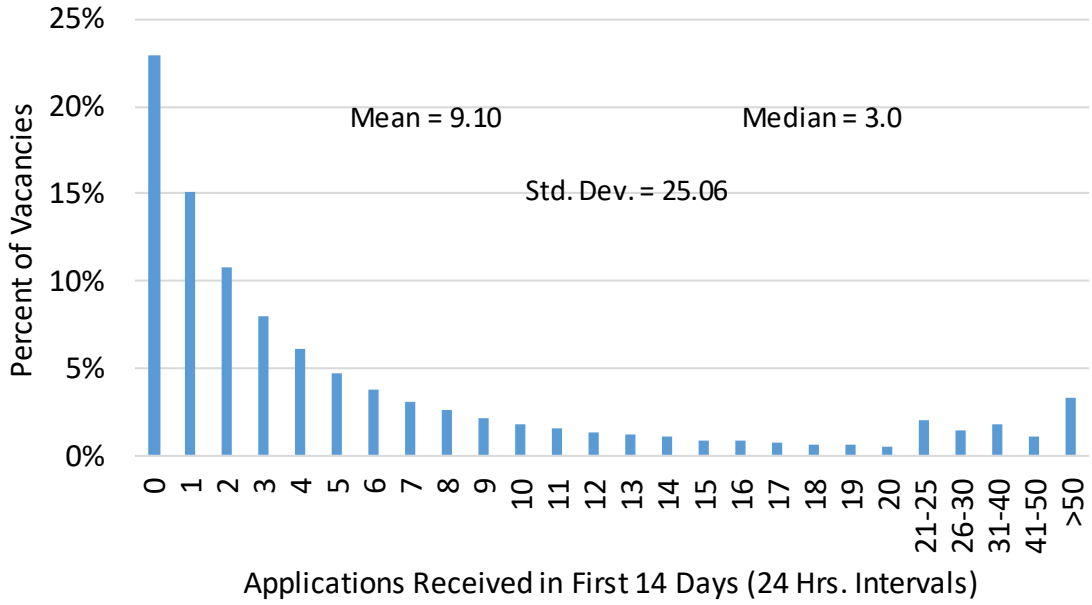
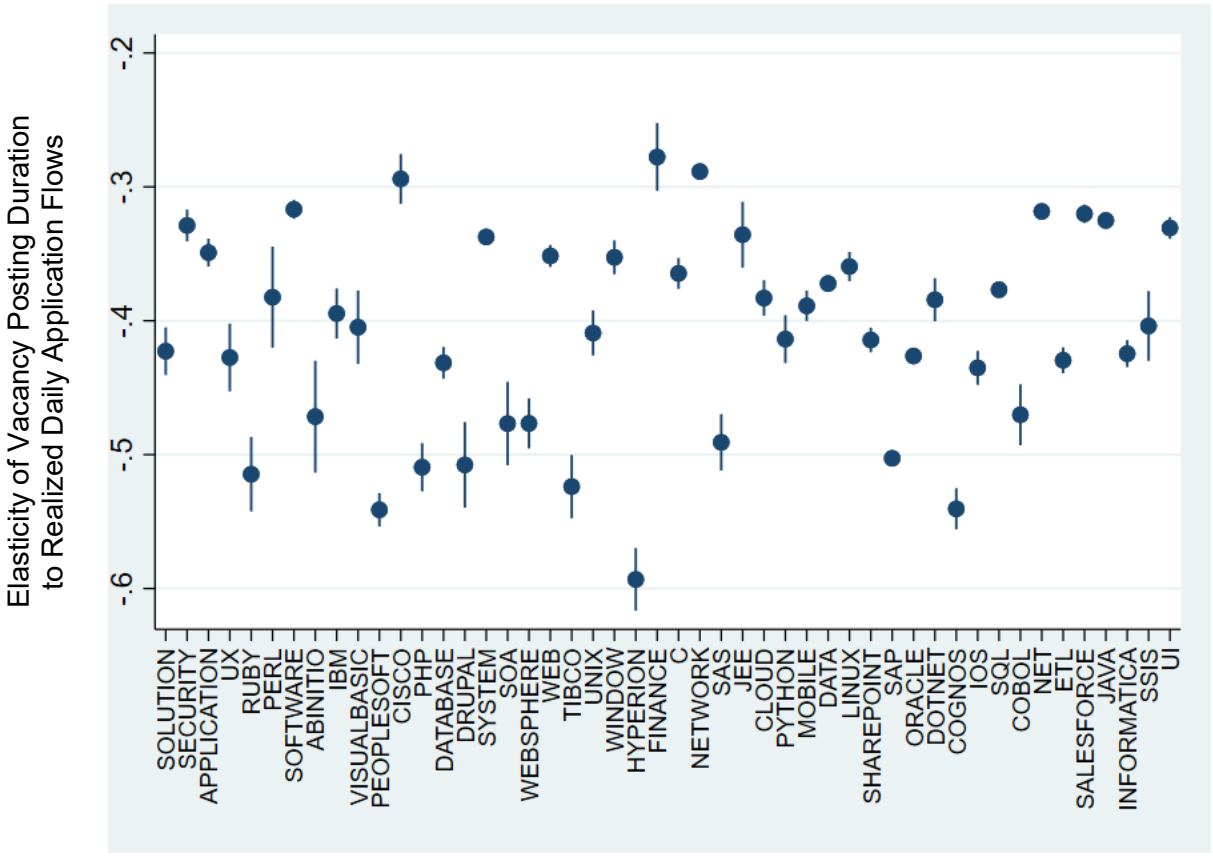
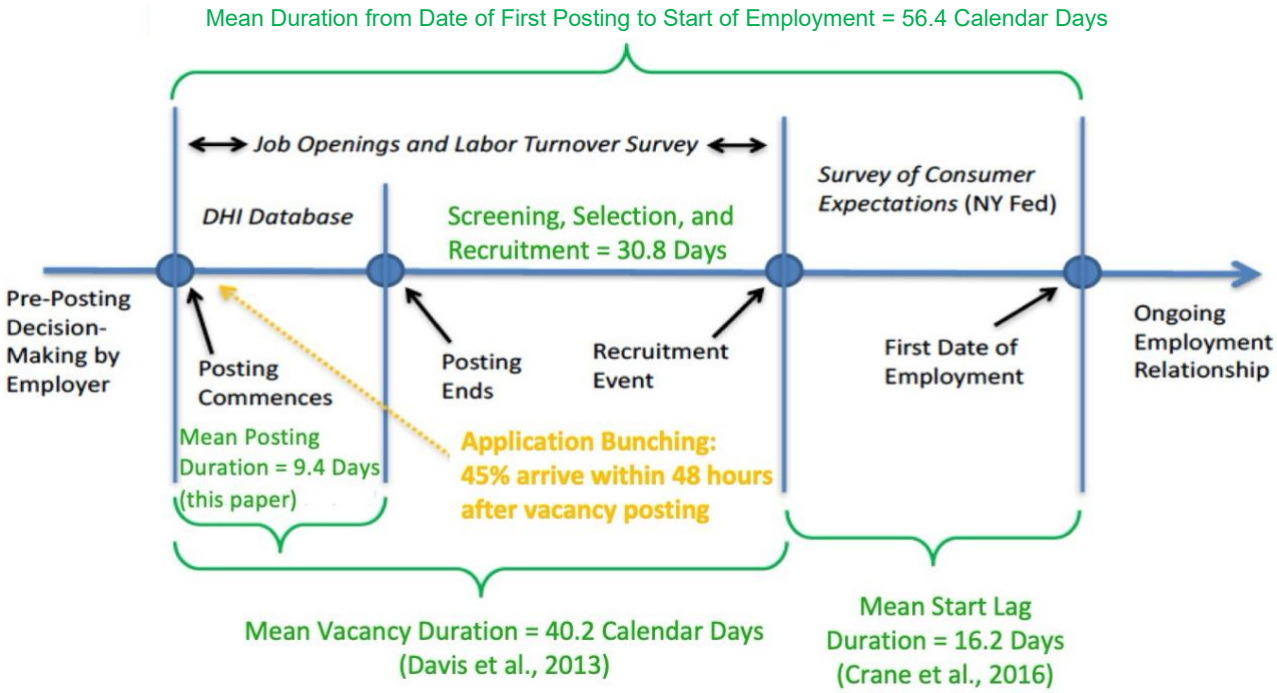


Figure 5. The Elasticity of Posting Duration with Respect to Application Flows by Skill Category, Conditional on Labor Market Tightness and Job Function



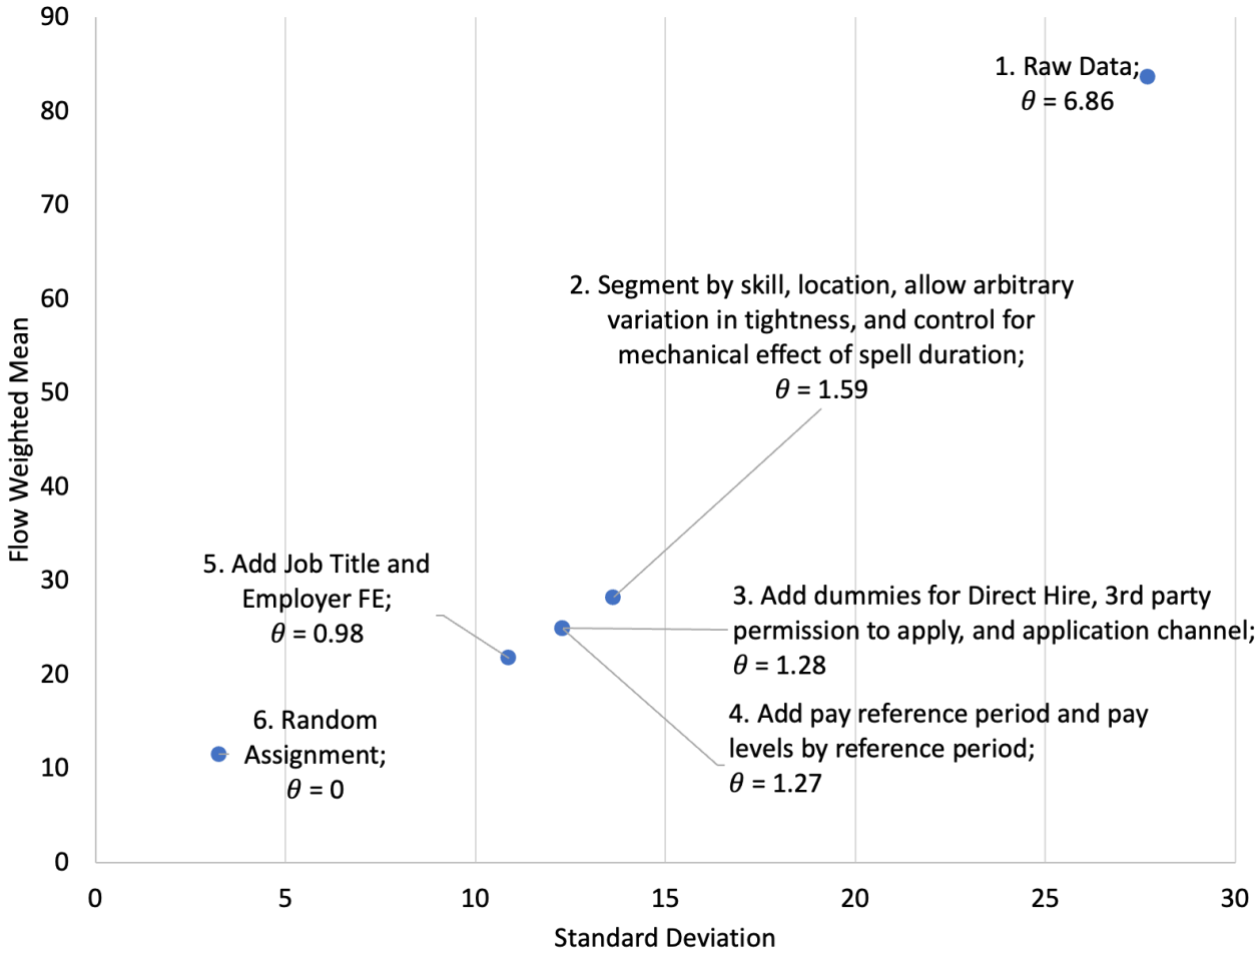
Notes: Each dot reports the estimated elasticity of posting duration with respect to daily application flows by skill category, controlling for job function fixed effects and skill-by-month fixed effects. We order skill categories from the lowest value of applications per posting on the left (SECURITY with 3.7 applications per posting) to the highest value on the right (User Interface with 24.8). The regression sample runs from January 2012 to December 2017 and covers all postings in the 48 Skill categories with at least 25 active postings in every calendar month. It excludes postings with first active date on or after 1 December 2017 to avoid incomplete spells.

Figure 6. Stages of the Hiring Process – A Quantitative Sketch



Notes: 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 using Job Openings and Labor Turnover Survey data and the methodology developed in Davis et al. (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. base 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 Figure 2 and Table 5. This figure is illustrative. Employers may begin recruitment activities before posting on Dice.com, and they may use other recruitment channels and platforms before or after a posting spell on Dice.com.

Figure 7. Selected Moments for the Distribution of Applications Over Postings: Raw Data, Random Assignment, and Fitted Binomial Models with Covariates



Notes: Point 1 reports selected moments in the raw data, and Point 6 reports the same moments under random assignment – i.e., for a binomial model that targets the mean of applications per posting. Points 2, 3, 4 and 5 report the estimated dispersion parameter (θ) for negative binomial models with covariates. The location of each point reflects the implied (conditional) values of the flow-weighted mean of applications per posting and the standard deviation of applications per posting. Table 9 reports coefficient estimates for selected covariates.

Appendix A: Additional Information about Data Processing

1. Long-Duration Postings

As remarked in Section II.4, one-quarter of the postings in the DHI Database are of the "long-duration" type. Because these long-duration postings typically pertain to multiple job openings – and even multiple employers in some cases – we set them aside in the analyses presented in the main text.[40] Nevertheless, there is potentially useful information in these long-duration postings. Thus, we bring them into some of the analyses reported below.

2. Out-of-Range and Repeat Applications

"Out-of-range" applications have a date-time stamp outside the interval defined by the posting's first and last active dates in the Activity File. Since postings should be visible to applicants only when active, an out-of-range application is one for which the date-time stamp on the application or the posting itself is misreported. We drop out-of-range applications from the sample. They account for 0.2% of all applications and occur for 0.6% of all postings.

Among the 125.3 million applications in our sample, 8.7% are repeats in the sense that a given applicant ID applies to a particular Job ID more than once. We include repeats in our analysis samples because we believe that employers are likely to regard them as distinct applications. In the case of third-party applications, an intermediary may submit applications on behalf of multiple job seekers using the same applicant ID. In the case of long-duration postings, a job seeker may apply to the same Job ID at different points in time. The case for excluding repeats is stronger for URL applications as these can arise if a job seeker clicks through to an external application system more than once to complete a previously initiated application. Rather than using different rules in these and other cases, we retain all repeats.

Table A.1 shows repeat applications as a percent of total applications by application channel and client type for standard and long-duration postings. As expected, the share of duplicate applications is higher for long duration postings, and for postings that redirect to an external URL to collect applications. Job postings from Direct Hire and Recruitment & Staffing Firms have similar shares of duplicate applications.

[40] 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 on Dice.com, according to DHI staff.

Table A.1. Total Applications and Repeats, January 2012 – December 2017

Panel A: Standard Postings					
	Email Applications		URL Applications		Total
	Direct Hire	Recruitment & Staffing Firms	Direct Hire	Recruitment & Staffing Firms	
Total Applications (Millions)	17.4	31.3	4.4	9.2	62.3
Repeats, percent of total	4.5	4.5	14.5	12.3	6.4

Panel B: Long-Duration Postings					
	Email Applications		URL Applications		Total
	Direct Hire	Recruitment & Staffing Firms	Direct Hire	Recruitment & Staffing Firms	
Total Applications (Millions)	17.0	29.5	9.0	7.4	63.0
Repeats, percent of total	9.0	8.8	17.2	16.7	11.0

Note: Repeats equal the difference between total applications and the sum of distinct applicant IDs.

Appendix B: Job Titles, Job Functions, and Skill Categories

In processing the text in the job-title field, we first delete punctuation and remove language about the company or work environment such as "Exciting," "Innovative," and "Urgent Need." We then replace acronyms and standardize common terms. These steps yield about 2 million distinct job titles, most of which involve few distinct postings. Parts of our analysis focus on common job titles with many postings. There are 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 common job titles account for over 93 percent of the Job IDs, Vacancy Posting IDs, and applications in the database. Appendix Table B.1 lists the most common job titles in our data. A full list of all 117,146 job titles with at least 250 distinct Job IDs (summed over both client types) is available here

Table B.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	BUSINESS SYSTEMS ANALYST	30,660

Notes: "Job ID Count" equals the number of distinct Job IDs (i.e., vacancy postings).

We also use the job-title text to group postings into Job Function and Skill categories. "Job Function" refers to our grouping of postings into 56 occupational categories such as "Programmer," "Developer," "Mechanical Engineer," "Consultant," and "Business Analyst." "Skills" refer to specific requirements mentioned in the job-title text. We consider 54 Skills such as "C," "SQL," "Java," "User Interface," and "Big Data." When a posting specifies multiple Job Functions (or Skills) that we cover, we use the first category mentioned in the job-title text. Our classification by job function covers 90 percent of postings while the skill categorization covers 55 percent of postings.

Table B.2 reports summary statistics for selected Skill Requirement categories. Job ID count refers to the sum of Job ID's in the Skill Requirement Category. Number of job titles is the sum of distinct job titles that include the Skill Requirement.

Table B.2. Selected Skill Requirement Categories in the DHI Database

Panel A. All Postings

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

Panel B. Standard Postings Only

Skill Requirement Category	*Job ID Count*	*Number of Distinct Job Titles in the Category*		*Ratio of Weighted to Unweighted Mean Daily Applications Per Posting*	*Average over Job Titles of the Ratio of Weighted to Unweighted Mean Daily Applications Per Posting*
		All	*With at least 100 postings*		
JAVA	312,933	198	50	6.2	5.4
SYSTEM	257,680	294	89	5.0	3.9
SOFTWARE	201,452	252	63	5.7	5.3
SAP	198,167	239	53	2.3	2.2
ORACLE	185,789	198	56	3.4	2.9
NETWORK	165,932	222	61	6.4	5.1
NET	163,889	190	38	6.1	5.5
DATA	138,621	268	58	3.7	3.1
APPLICATION	106,660	243	62	4.3	3.9
WEB	100,215	209	45	7.7	6.8
SECURITY	100,205	235	52	4.6	3.6
SQL	105,079	177	36	3.6	3.5

Notes: "Job ID Count" equals the number of distinct Job IDs in the indicated Skill Requirements Category. "Number of Distinct Job Titles" is the number of Job Titles represented among the Job IDs grouped into the indicated category. For the rightmost two columns, we first compute daily applications per posting as applications received divided by posting duration (days elapsed from the first to last date-time on which the posting was in active status). To obtain entries for the second column from the right, we compute the ratio of (a) the flow-weighted mean of daily applications per posting for postings in the indicated category to (b) the unweighted mean of daily applications per posting in the category. To obtain entries in the rightmost column, we compute the ratio of (a) the flow-weighted mean of daily applications per posting at the Job Title level to (b) the unweighted mean of daily applications per posting at the same level. Finally, we compute the simple mean of these ratios over Job Titles represented in the Skill category and report it in the rightmost column.

Appendix C: Additional Empirical Results

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

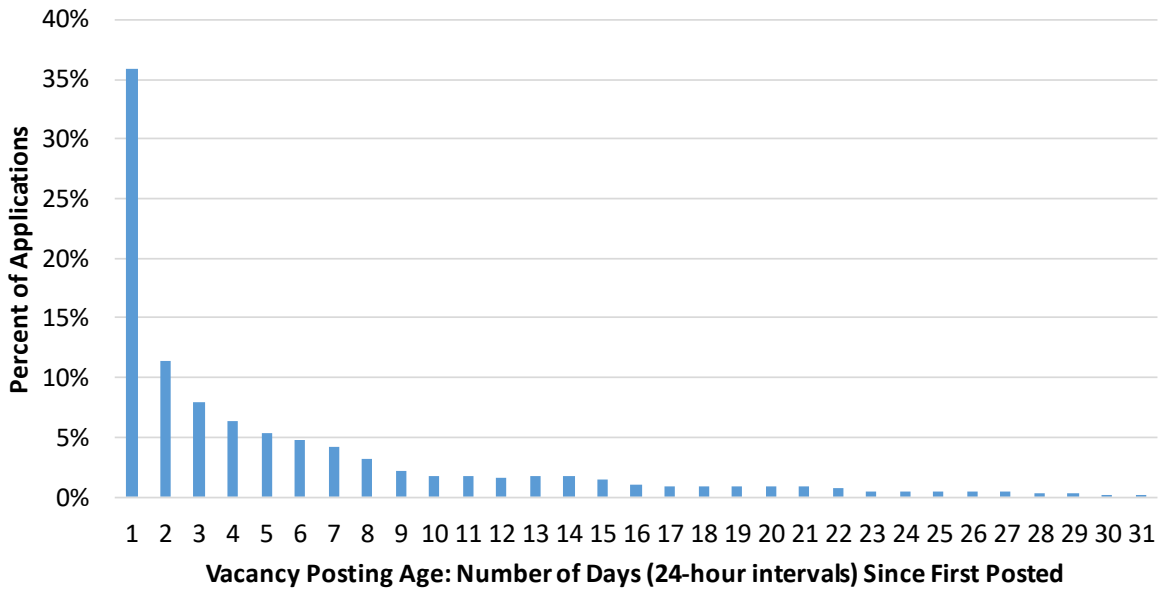
			Percentile				
	No. of Standard Postings	Mean	10	25	50	75	90
All Job Titles with at Least 100 Standard Postings	5,157,666	9.44	1.00	2.93	7.00	14.02	22.71
Employer Type (ownership)							
Privately Held Companies	4,744,376	9.35	1.00	2.82	6.92	14.03	22.56
Publicly Listed Companies	258,737	11.15	5.00	7.00	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. In constructing this table, we drop observations with first posting date on or after December 1, 2017 to avoid incomplete spells.

Figure C.1 displays the distribution of applications by posting age separately for Direct Hire Clients and Recruitment & Staffing Firms.

Figure C.1. The Distribution of Applications by Vacancy Posting Age, Standard Postings, January 2012 to December 2017

A. Direct Hire Clients



B. Recruiting and Staffing Firms

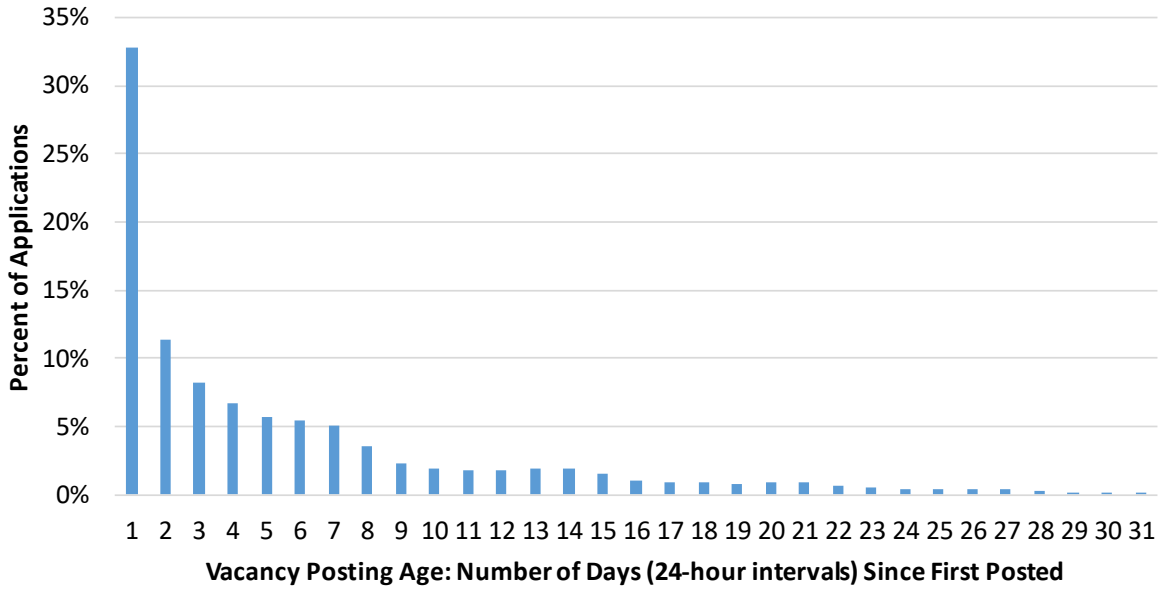
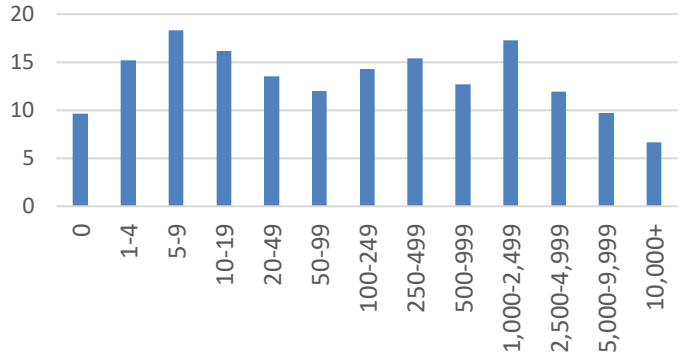


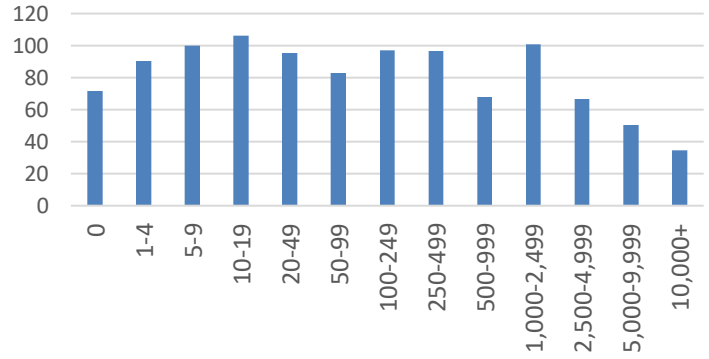
Figure C.2 shows weighted and unweighted mean applications per posting by employer size for Direct Hire clients. Perhaps surprisingly, there is no strong, simple relationship between employer size and the size of applicant pools. The largest employers draw the smallest applicant pools, while employers with 1,000 to 2,500 employees draw relatively large pools. 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. From the applicant perspective (Panel B in Figure C.2), competition is similar at firms with 10 to 19 employees, 100 to 500 employees, and those with 1,000 to 2,499 employees. In unreported results, controlling for differences in the mix of job titles does not greatly alter the relationship between employer size and mean applications per posting.

Figure C.2. Mean Applications per Vacancy by Employer Size, January 2012 to December 2017

Panel A. Direct Hires, Standard Postings, Job Titles with 100+ Standard Postings, Equal Weights

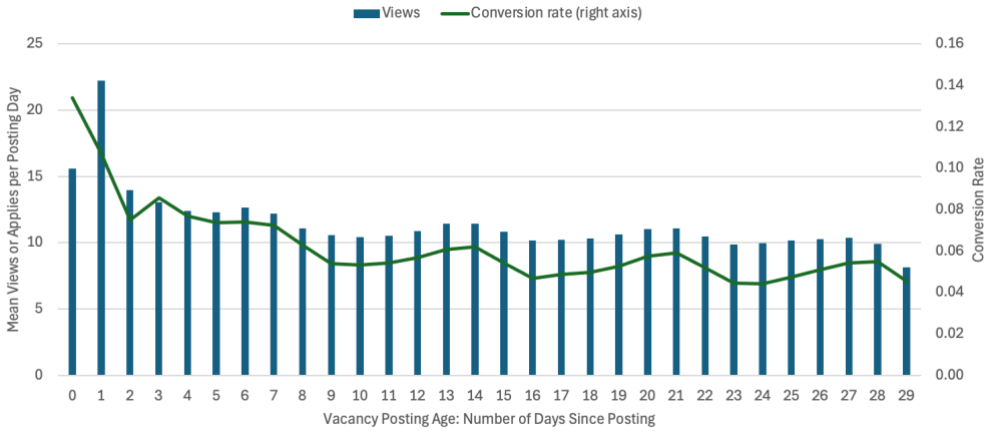


Panel B. Direct Hires, Standard Postings, Job Titles with 100+ Standard Postings, Weighted by Application Flows



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 instead of focusing on job titles with at least 100 job postings.

Figure C.3: Mean Daily Views and Conversion Rates by Posting Age, Standard Postings



Notes: 0 on the x-axis indicates the day of first posting. 4.7% of posting days appear to have zero views but positive applications in the data. These are kept as recorded. We calculate the conversion rate as the sum of applications across all postings in the age bin divided by the sum of views across all postings in the same age bin.

Table C.2 below reports estimated elasticities of posting durations with respect to slack for labor markets defined by job functions, MSAs, and MSA-skill cells. These tables and the underlying regression models parallel those in Table 7 in the main text.

Table C.2. How Posting Durations Vary with Slack, Alternative Market Definitions

Panel A. Labor Markets Defined by 34 Job Function Categories

Panel A1. Main Regressions

Dependent Variable: ln(Mean Duration of Postings in Job Function *j* in Month *t*)

	(1)	(2)	(3)	(4)	(5)	(6)
ln(Job Seekers/Postings)	0.024***	0.047***			0.112***	0.152***
	(0.006)	(0.007)			(0.014)	(0.013)
ln(Applications/Postings)			0.002	0.008*	-0.066***	-0.087***
			(0.004)	(0.005)	(0.009)	(0.009)
Constant	2.36***	2.29***	2.39***	2.32***	2.33***	2.27***
	(0.008)	(0.015)	(0.010)	(0.016)	(0.012)	(0.017)
Observations	2,414	2,414	2,414	2,414	2,414	2,414
R-squared	0.65	0.74	0.65	0.73	0.66	0.74
Within R-squared	0.01	0.02	0.00001	0.001	0.03	0.06
Fixed Effects	JF	JF & Time	JF	JF & Time	JF	JF & Time

Panel A2. Auxiliary Regressions

Dependent Variable: ln(Applications Per Posting in Job Function Category *j* in Month *t*)

	Elasticity of Applications Per Posting with Respect to Job Seekers Per Posting	Within R-squared
Controlling for Job Function Fixed Effects	1.34 (0.013)	0.81
Controlling for JF and Time Fixed Effects	1.21 (0.014)	0.76

Panel A3. Selected Summary Statistics

	Log Posting Durations	Log Job Seekers Per Posting	Log Applications Per Posting
Standard Deviation Across Job Function-Month Cells	0.15	0.49	0.71
Average Standard Deviation over Time within JF Categories	0.08	0.27	0.40

Table C.2. (continued)
Panel B. Labor Markets Defined by 53 Metropolitan Statistical Areas (MSAs)

Panel B1. Main Regressions
Dependent Variable: ln(Mean Duration of Postings in MSA *j* in Month *t*)

	(1)	(2)	(3)	(4)	(5)	(6)
ln(Job Seekers/Postings)	0.005*	0.149***			-0.068***	0.050***
	(0.003)	(0.006)			(0.014)	(0.014)
ln(Applications/Postings)			0.007***	0.135***	0.067***	0.095***
			(0.003)	(0.005)	(0.013)	(0.012)
Constant	2.11***	1.99***	2.10***	1.95***	2.08***	1.96***
	(0.015)	(0.018)	(0.015)	(0.018)	(0.016)	(0.018)
Observations	5,964	5,964	5,964	5,964	5,964	5,964
R-squared	0.42	0.55	0.42	0.55	0.42	0.55
Within R-squared	0.001	0.11	0.001	0.12	0.01	0.12
Fixed Effects	MSA	MSA & Time	MSA	MSA & Time	MSA	MSA & Time

Panel B2. Auxiliary Regressions
Dependent Variable: ln(Applications Per Posting in MSA Category *j* in Month *t*)

	Elasticity of Applications Per Posting with Respect to Job Seekers Per Posting	Within R-squared
Controlling for MSA Fixed Effects	1.10 (0.003)	0.96
Controlling for MSA and Time Fixed Effects	1.04 (0.006)	0.83

Panel B3. Selected Summary Statistics

	Log Posting Durations	Log Job Seekers Per Posting	Log Applications Per Posting
Standard Deviation Across MSA-Month Cells	0.16	0.57	0.65
Average Standard Deviation over Time within MSA Categories	0.11	0.51	0.57

Table C.2. (continued)
Panel C. Labor Markets Defined by 113 MSA-Skill Categories

Panel C1: Main Regressions
Dependent Variable: ln(Mean Duration of Postings in MSA×Skill *j* in Month *t*)

	(1)	(2)	(3)	(4)	(5)	(6)
ln(Job Seekers/Postings)	0.025***	0.112***			0.057***	0.106***
	(0.003)	(0.004)			(0.011)	(0.011)
ln(Applications/Postings)			0.020***	0.092***	-0.029***	0.005
			(0.002)	(0.004)	(0.010)	(0.010)
Constant	2.16***	2.11***	2.15**	2.06***	2.182***	2.102***
	(0.017)	(0.019)	(0.017)	(0.020)	(0.018)	(0.020)
Observations	8,023	8,023	8,023	8,023	8,023	8,023
R-squared	0.54	0.62	0.53	0.61	0.54	0.62
Within R-squared	0.01	0.08	0.01	0.07	0.01	0.08
Fixed Effects	MSA-Skill	MSA-Skill & Time	MSA-Skill	MSA-Skill & Time	MSA-Skill	MSA-Skill & Time

Panel C2. Auxiliary Regressions
Dependent Variable: ln(Applications Per Posting in MSA-Skill Category *j* in Month *t*)

	Elasticity of Applications Per Posting with Respect to Job Seekers Per Posting	Within R-squared
Controlling for MSA-Skill Fixed Effects	1.10 (0.003)	0.95
Controlling for MSA-Skill and Time Fixed Effects	1.05 (0.005)	0.86

Panel C3. Selected Summary Statistics

	Log Posting Durations	Log Job Seekers Per Posting	Log Applications Per Posting
Standard Deviation Across MSA-Skill-Month Cells	0.16	0.57	0.65
Average Standard Deviation over Time within MSA-Skill Categories	0.11	0.51	0.57

Notes: The sample includes standard job postings with first active dates between January 2012 and November 2017. We group postings into skill categories (job functions) by the first skill (job function) mentioned in the posting's extended job title, and we group them into MSAs based on the location of the job. We group postings into months based on the month in which the posting first became active. See the notes to Table 7 in the main text for additional information.

Table C.3. Selected Statistics and Negative Binomial Models with Covariates by Skill Category, Categories with the Most Postings

Unit of Analysis: Completed Posting Spell
Dependent Variable: Number of Applications to the Posting

			Applications per Posting		*Maximum Likelihood Estimates of the Dispersion Parameter, θ*								
Skill Category	Number of Postings	Percent of Postings with No Applications	Simple Mean	St. Dev.	No Co-Variates	+ Spell Dur-ation	+ Monthly Time Effects (71)	+ State FEs (54)	+ Pay-period indicators (3) + ln(real pay) X pay-period indicators (3)	+ 3rd party Okay (1)	+ Job Titles FEs (1,626)	+ Employer Size Fixed Effects (14)	+ Employer FEs (4,115)
APPLICATION	94,721	26%	5.2	10.9	2.03	1.71	1.61	1.51	1.51	1.39	1.21	1.19	0.93
DATA	119,639	16%	9.9	19.2	1.85	1.71	1.38	1.34	1.33	1.09	0.92	0.91	0.74
DOTNET	16,453	16%	11.6	25.8	2.17	2.11	1.36	1.27	1.24	0.86	0.82	0.79	0.50
JAVA	281,798	17%	19.5	58.8	2.89	2.83	1.81	1.78	1.75	1.24	1.14	1.13	0.93
NETWORK	144,856	23%	8.4	21.9	2.47	2.35	1.81	1.75	1.74	1.32	1.22	1.17	0.93
ORACLE	169,488	14%	11.5	22.9	1.86	1.71	1.29	1.26	1.25	1.13	0.93	0.93	0.78
SAP	177,400	11%	11.3	17.3	1.44	1.33	1.12	1.11	1.10	1.02	0.98	0.97	0.81
SECURITY	84,922	30%	4.1	8.5	1.95	1.58	1.43	1.37	1.37	1.16	0.99	0.96	0.71
SOFTWARE	176,673	26%	6.1	14.9	2.26	2.02	1.88	1.80	1.79	1.60	1.36	1.33	1.05
SQL	95,346	12%	17.1	36.9	2.02	1.92	1.45	1.41	1.36	0.95	0.90	0.88	0.70
SYSTEM	227,179	23%	6.9	14.4	2.05	1.87	1.65	1.59	1.58	1.32	1.10	1.08	0.87
WEB	91,729	27%	7.4	22.4	2.83	2.74	1.96	1.90	1.88	1.52	1.36	1.33	1.04

Notes: The first column reports the skill category, and the next four columns report selected statistics for the indicated category. The remaining columns report maximum likelihood estimates of the dispersion parameter, θ, for increasingly expansive versions of the NB model (2) and (3) set forth in Section IV.7. See the notes to Table 9 for information about the sample. When the posting-level observation has missing data, we introduce a missing category and treat it as one of the classification levels. For example, our State Fixed Effects cover all 50 individual states, Puerto Rico, DC, other, and missing. The "other" category covers "USA," "Nationwide," "100% travel" and some combination of letters and numerals.

Appendix D: Analytical Relationship between the Simple Mean Number of Applications per Vacancy Posting and the Flow-Weighted Mean

Proposition: Consider *v* vacancy postings, where v_n of them attract $n = 0, 1, 2, \ldots, n^{max}$ applications. Let *M* and σ^2 denote the unweighted mean and variance of the distribution of application flows over the *v* postings, and let M^W denote the flow-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 the v_n postings with *n* applications apiece. The probability function of postings over the number of applications 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, write the flow-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.

Consider the case in which applications flow to postings in a completely random manner. Specifically, there is a uniform probability that any given application flows to any given posting. 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 proposition that $M^W = (a/v) + 1 - (1/v)$ for the binomial case, and that M^W goes to $(a/v) + 1$ for a large number of vacancy postings.

Appendix E: Additional Remarks on the Literature

Posting Durations on Other Online Job Boards

Meaningful comparisons of posting durations on Dice.com to posting durations on other online job boards are challenging due to measurement issues and differences in pricing models across platforms. Consider two cases. First, Marinescu and Wolthoff (2020) 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 period (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 covered a 60-day period. They include postings that pertain to multiple job openings, which typically have much longer durations.

Other Evidence of How Intermediaries Operate on Labor Market Matching Platforms

Stanton and Thomas (2016) consider the role of intermediaries on oDesk.com (later known as Upwork), an online platform for spot contracts in remotely supplied labor services. Most of the contracts transacted on oDesk involve employers in high-income countries that retain the services of workers in low-income countries. Interestingly, and to "the surprise of oDesk's management, more than 1,100 small autonomous outsourcing agencies entered the oDesk market within a few years. These agencies operate within the oDesk platform, but contracting, monitoring, and work direction still take place between employers and individual workers." These agencies provide screening and vetting services that overlap with the services that Recruitment and Staffing firms provide on Dice.com.

Horton (2017) studies a field experiment in which oDesk offered algorithm-based recommendations to employers. The recommendations raised the success rate in forming matches by 20 percent in technical job openings, with no apparent evidence that other matches were crowded out.

How Non-Sequential Search Can Affect the Equilibrium Wage Structure and Job Types

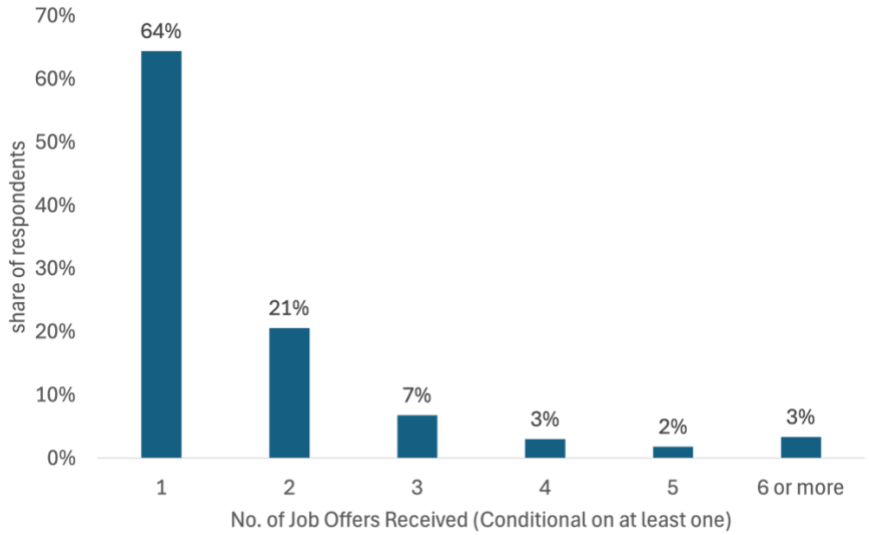
Non-sequential search injects distinct forces into the determination of equilibrium outcomes. To see this point, suppose job seekers can submit multiple applications while awaiting news about callbacks and offers, can take at most one job, and run the risk of no offers. This search problem has aspects of portfolio choice in that the number *and* mix of vacancies to which the job seeker applies affect his expected payoff. As Chade and Smith (2006) prove, it is not then generally sufficient to rank order vacancies by expected payoffs and then optimize over the number of applications. Instead, when jobs differ enough in attractiveness and offer probabilities, and if costs per application are not too high, the optimal non-sequential strategy is to apply to a mix of highly attractive and not-so-attractive jobs while foregoing jobs in the middle (Chade and Smith, 2006). Galenianos and Kircher (2009) integrate this portfolio choice perspective into an equilibrium model of directed search. In their model, job seeker appetites for both "risky" job openings (high wage, low offer probability) and "safe" ones (low wage, high offer probability) support equilibrium wage dispersion with homogenous agents. The number of simultaneous applications per job seeker determines the extent of wage dispersion and, hence, the types of jobs that emerge in equilibrium.

How the Best Offer Wage Varies with the Number of Job Offers

As remarked in Section IV.3, job seekers strengthen their bargaining positions when a non-sequential search strategy yields multiple job options. This claim has a common-sense quality to it, but we are unaware of previous studies that empirically explore how a jobseeker's best offer wage varies with the number of offers received.

We investigate this matter using data from the New York Federal Reserve Board's Survey of Consumer Expectations Job Search Module. The following histogram shows the distribution of the number of offer wages received by job seekers in this dataset.

Figure E.1. Histogram of Job Offers Received in the Last Six Months



As reported in Table E.1 below, we find a positive, sizable, and statistically significant relation between the number of offers a job seeker receives and the best offer wage. According to the linear specification in column (1), a 1% increase in the number of offers is associated with a 12.6% increase in the hourly wage of the best offer. (It is 14.3% when restricting attention to persons with at least one offer in the previous four weeks.) Column (2) in the table reports results for a nonparametric specification. While the modest sample size precludes precise inference, the best offer wage rises with the number of offers. For example, the best offer wage is, on average, 36% higher for a jobseeker with six offers in the past six months as compared to someone with only one offer in the past six months. Columns (3) and (4) add a battery of controls for education, sex, and age to the linear and nonparametric specifications. The results are similar to columns (1) and (2).

Table E.1. How the Best Wage Offers Varies with the Number of Offers Received

Dependent Variable: ln(hourly wage at best offer)				
	(1)	(2)	(3)	(4)
ln(number of job offers)	0.126***		0.111***	
	(0.044)		(0.038)	
number of offers				
2		0.049		0.045
		(0.060)		(0.048)
3		0.145		0.145
		(0.096)		(0.092)
4		0.135		0.214**
		(0.131)		(0.091)
5		0.163		0.158
		(0.135)		(0.103)
6 or more		0.357**		0.258*
		(0.165)		(0.160)
Constant	2.802***	2.810***	1.230***	1.226***
	(0.032)	(0.035)	(0.287)	(0.290)
Observations	1,690	1,690	1,688	1,688
R-squared	0.0109	0.0106	0.2541	0.2540
Controls	No	No	Education category, gender, age, age^2	

Note: Our analysis sample contains all jobseekers in the New York Fed's Survey of Consumer Expectations Job Search Module with at least one job offer in the past six month.

Other Studies that Relate the Behavior of Job Seekers to Directed Search Models

A few other empirical studies relate the behavior of job seekers to directed search models. Banfi and Villena-Roldán (2019) find support for two core implications in data from an online Chilean job board: First, postings that offer higher wages attract more applicants. Second, the experience and education requirements specified in job ads correlate positively with the average qualifications of the applicant pools. Belot et al. (2022) find that higher offer wages attract more applicants in a field experiment that features random variation in offer wages across otherwise identical vacancy postings. Marinescu and Wolthoff (2020) find that applicant numbers rise with posted wages in U.S. data from Careerbuilder.com in 2011, but only after conditioning on job titles. In contrast, Faberman and Menzio (2018) find that applicant numbers fall with posted wages in U.S. survey data from the early 1980s. They develop a model of directed search with two-sided heterogeneity to rationalize this finding.

As a theoretical matter, we note that the decision of *whether* to specify the wage in the vacancy posting can also play a role in directing the flow of applications. Michelacci and Suarez (2006) consider a search model where worker productivity is observable but not verifiable. They identify circumstances in which employers choose to forego the benefits of wage posting in job ads to attract applicants of higher quality.

Posting Durations, Market Tightness, and Imperfect Classification of Labor Markets

Posting durations on Dice.com exhibit almost no tendency to lengthen with market tightness, as we show in Table 7 and discuss in Section III.7. One might be concerned that this finding is simply an artifact of measurement error. That is, imperfect classification of labor markets could attenuate the estimated relationship between posting durations and tightness. We now take up this concern.

Many studies find strong evidence of equilibrium relationships between tightness and various outcome variables for labor markets defined by industry, occupation, area, or their cross products. Here are a few examples:

- Using JOLTS data, Davis et al. (2012) find a strong relationship between the change in vacancy durations over time and the change in market tightness at the industry level. See their Figure 3.a, which shows that industry-level vacancy durations rise strongly with industry-level tightness. Equivalently, in their presentation, industry-level job-filling rates fall strongly with industry-level tightness.
- Carrillo-Tudela et al. (2023) find that vacancy yields (i.e., vacancy-filling rates) vary systematically with tightness across 36 region-by-skill markets in Germany. See their Figure 5.
- Azar et al. (2022) find that real wages rise with tightness across markets defined at the level of six-digit occupation classifications crossed with commuting zones. This is a highly granular classification of labor markets, which increases the scope for imperfectly defined, overly narrow markets to attenuate the estimated relationship between tightness and market-level outcomes. Still, the authors find strong statistical evidence that tighter markets have higher wages, conditional on many other covariates.
- Bilal (2023) finds that the job-finding rate of unemployed persons covaries positively with labor market tightness across commuting zones in France and in the United States. See his Figures II.B and VI.B.
- The design of labor mismatch measures and the empirical implementation of mismatch models rests on the premise that standard industry classifications, occupation classifications, local market areas (MSAs, CZs, etc.) or some combination of all three capture relevant labor markets with sufficient accuracy to allow for meaningful empirical analysis. Şahin et al. (AER) is a prominent example.

In light of the many studies that find meaningful relationships between measured market-level variation in tightness and other market-level outcomes for markets defined by skill, industry, occupation, MSAs or their cross products, we see measurement error as an untenable explanation for our finding that posting durations do not covary with market tightness for markets defined by skills, job functions, MSAs, and the cross product of MSAs and skills.

Appendix F: How Posting Durations Relate to Idiosyncratic Application Flows – An Alternative Investigation

We now investigate how posting durations relate to idiosyncratic application flows at a daily frequency. This investigation complements the analysis in Section III.6 in the main text. While the empirical specification considered here is more complex than the one that underpins Figure 5 and less amenable to a compact presentation of results, it has two appealing features. First, it lets us explore how the relationship between idiosyncratic application flows and posting termination varies with posting duration. Second, the specification admits posting-level fixed effects, allowing us to control for unobserved attributes of the posting and the circumstances of its posting.

Consider posting *i* on each day *d* that it's active (i.e., open to applications) on the Dice platform. Let $\bar{d}_i$ be the last day posting *i* is active, and let $Y_{i,d}$ be an indicator variable equal to 1 if $d = \bar{d}_i$ and zero otherwise. We estimate the following specification:

$$Y_{i,d} = \sum_{d=0}^{\bar{d}_i} \beta^d I[d_i = d] \times a_{i,d} + day_{d,i} + job_i + \eta_{t(i)\times s(i)} + \epsilon_{i,d,t}$$

where $a_{i,d}$ is the inverse hyperbolic sine of daily applications received by posting *i* on day d. $day_{i,d}$ is a set of fixed effects for how many days posting *i* has been active when observed on day *d,* job_i is a fixed effect for posting *i*, $t(i)$ denotes the month when job *i* is first active, and $s(i)$ denotes the posting's skill category. Thus, $\eta_{t(i)\times s(i)}$ is a set of skill-time fixed effects, that capture market tightness and any other forces that vary over time at the market level. The $day_{i,d}$ effects are nonparametric controls for posting duration to date.

Figure F1 below plots the estimated β^d coefficients. The results are consistent with Figure 5: An increase in the idiosyncratic flow of applications on a given day raises the probability that the posting exits – i.e., is no longer active the next day. The magnitude of this effect falls slowly with duration, except on day 0 and on days 7, 14 and 21. (We interpret this latter pattern as evidence that many postings terminate after one, two or three weeks irrespective of applicant numbers.) A one percent idiosyncratic increase in the number of applicants on day 10, for example, raises the probability that the posting is no longer active on the next day by about 4.2 percentage points. The only feature of Figure F1 that departs from the prevailing pattern (and the chief message of Figure 5) is the negative value for $\beta^{d=0}$. This negative value says that an extra idiosyncratic flow of applicants on Day 0 (the first day the posting is active) lowers the probability that Day 0 is the last active day for the posting.

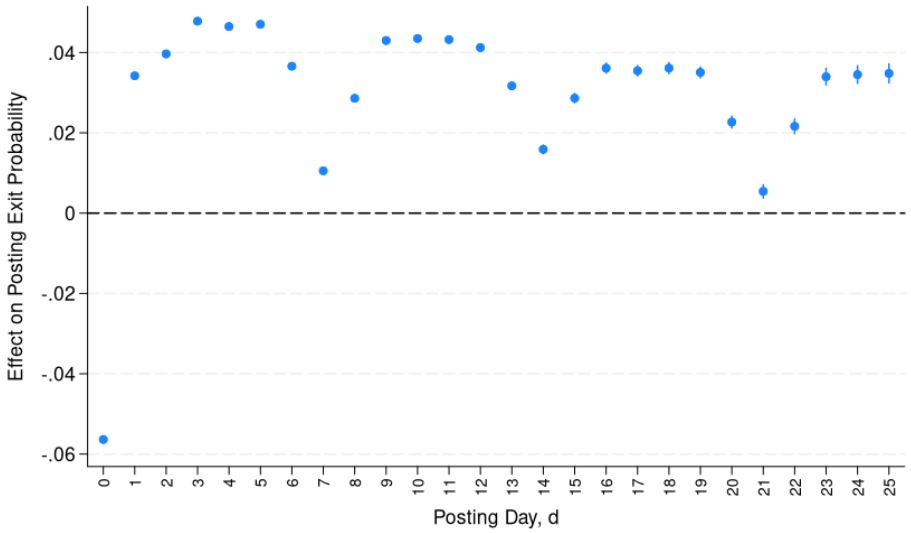
The magnitude of the estimates in Figure F1 and Figure 5 are consistent. If we approximate the exit process for a posting with a constant daily hazard (h), then mean duration is $E[d] = 1/h$ and $\ln(E[d]) = -\ln(h)$ and the duration elasticity with respect to applications evaluated at the mean duration is:

$$\frac{\partial lnE[d]}{\partial \ln(a)} = -\frac{1}{h}\frac{\partial h}{\partial \ln(a)} = -E[d] \times \beta$$

Therefore, the duration elasticity implied by estimates in Figure F1 is −0.04 × 9.4 = −0.38, directly comparable to the elasticity reported in the main text from Figure 5 of -0.4.

In sum, Figure F1 confirms the main message of Figure 5. However, Figure 5 is more compact in showing that our main result also holds across skill categories. We also think Figure 5 is easier for the reader to digest. For these reasons, we highlight Figure 5 in the body of the paper.

Figure F.1. Estimated Effects of Idiosyncratic Application Flows on the Posting Exit Probability as a Function of Duration Conditional on Posting-Level Fixed Effects, Posting Duration to Date, and Market-Level Tightness



Note: Each dot shows the point estimate for the indicated β^d. The whiskers show 95% confidence intervals.

Additional References for the Appendices

Azar, Jose, Ioana Marinescu and Marshall Steinbaum, 2022. "Labor Market Concentration," *Journal of Human Resources*, 57 (S), S167-S199.

Belot, Michèle, Philipp Kircher, and Paul Muller, 2022. "How Wage Announcements Affect Job Search – A Field Experiment," *American Economic Journal: Macroeconomics,* 14, no. 4, 1-67.

Bilal, Adrien, 2023. "The Geography of Unemployment," *Quarterly Journal of Economics*, 138, no. 3 (August), 1507-1576.

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

Carrillo-Tudela, Carlos, Hermann Gartner and Leo Kaas, 2023. "Recruitment Policies, Job-Filling Rates, and Matching Efficiency," *Journal of the European Economic Association*, 21, no. 6 (December), 2413-2459.

Chade, Hector and Lones Smith, 2006. "Simultaneous Search," *Econometrica,* 74, no. 5 (September), 1293-1307.

Faberman, R. Jason and Guido Meznio, 208. "Evidence on the Relationship between Recruiting and the Starting Wage," *Labour Economics,* 50 (March), 67-79.

Horton, John, 2017. "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment," *Journal of Labor Markets,* 35, no. 2.

Micehelacci, Claudio and Javier Suarez, 2006. "Incomplete Wage Posting," *Journal of Political Economy,* 114, no. 6 (December).

Şahin, Aşegül, Joseph Tong, Giorgio Topa and Giovanni L. Violante, 2014. "Mismatch Unemployment," *American Economic Review*, 104, no. 11 (November), 3529-64.