Deconstructing Job Search Behavior*

Stefano Banfi
Fiscalía Nacional Económica, Chile

Sekyu Choi
University of Bristol

Benjamín Villena-Roldán
Universidad Andres Bello, LM^{2}C^2, and MIPP.

May 20, 2023

Abstract

We use online job board data to document novel facts regarding unemployment and on-the-job searches. We define a relevant set of job ads using the bipartite network between ads and applicants and show how applications are affected by demographics such as gender, age, and marital status as well as timing variables like unemployment duration, job tenure, and cyclical conditions. We also study the selective margin, i.e., how the posted-expected wage gap and the applicant fit into ad requirements such as education, experience, location, and occupation affect applications. This evidence can help discipline current and future search theories.

Keywords: Online job search, Applications, Unemployment, On-the-job search, Networks.

JEL Codes: E24, J40, J64

*Email: sbanfi@fne.gob.cl, sekyu.choi@bristol.ac.uk and benjamin@benjaminvillena.com . We thank Guido Menzio, Brenda Samaniego de la Parra, Jan Eeckhout, Shouyong Shi, Yongsung Chang, Francois Langot, Bariş Kaymak, and seminar participants at Cardiff University, University of Manchester, Diego Portales University, the 2016 Midwest Macro Meetings, the 2016 LACEA-LAMES Workshop in Macroeconomic, the Search & Matching workshop at the University of Chile, 2017 SAM Meetings in Barcelona, Spain, 2017 SECHI Meeting, 2017 SM3 Workshop, and Universidad de Los Andes Bogotá, Université de Montréal, Universidad de Chile, the IDSC of IZA Workshop: Matching Workers and Jobs Online, and the 2022 RIDGE Labor Workshop for insightful discussion and comments. Villena-Roldán thanks for financial support the FONDECYT grant project 1191888; ANID Nucleo Milenio Labor Market Mismatch - Causes and Consequences, LM^{2}C^2 (NCS2022_045); and the Institute for Research in Market Imperfections and Public Policy, MIPP (ICS13_002 ANID). The data underlying this article were provided by www.trabajando.com through a legal agreement. Data will be shared on request to the corresponding author with the permission of the legal representative of www.trabajando.com. We are indebted to Ignacio Brunner, Álvaro Vargas and Ramón Rodríguez for valuable practical insights regarding the data structure and labor market behavior. All errors are ours.

1 Introduction

What kind of jobs workers look for and how much effort they exert are critical for labor market outcomes. Job search determines wages and job allocations in the economy to a large extent. As applications are tentative allocations, they are of prime importance to understand ex post outcomes such as the unemployment rate, mobility, efficiency, and income inequality. Thus, we use data from the job posting website `www.trabajando.com` to deconstruct behavior into two margins and quantify their impact on application decisions: First, an *intensive margin*, where we focus on characteristics of applicants or job postings, as studied in Faberman and Kudlyak (2019); Gomme and Lkhagvasuren (2015); Mukoyama, Patterson, and Şahin (2018). Second, a *selective margin*, where the focus is on coincidences between characteristics of the worker and the job. To the best of our knowledge, we are the first to systematically study both margins jointly and characterize the job search behavior of both unemployed and employed jobseekers.

Realistically, every time a worker visits `www.trabajando.com` use filters to define the scope of the search to a subset of relevant ads among thousands available, what the literature in industrial organization and marketing refers to as a *consideration set* (Van Nierop, Bronnenberg, Paap, Wedel, and Franses, 2010; Honka, Hortaçsu, and Wildenbeest, 2019; Abaluck and Adams-Prassl, 2021). Since clicks on job ads are unobserved to us, we use the *network* formed by job seekers linked through common job applications to construct the choice set of each applicant as the list of all ads applied by seekers linked with her. Our methodology not only overcomes the problem of only observing actual choices in the data but also has two important advantages over alternatives in the literature:[1] First, it relies on applicant-revealed preferences over a probably large number of observable and unobservable (to the econometrician) job characteristics to define similarities among jobs instead of a few arbitrary dimensions such as location, occupation, or industry. Using arbitrary choice sets, one would inevitably ignore attempted mobility across segments, a likely important issue as shown in Carrillo-Tudela and Visschers (2023). Second, the network approach allows us to define individual choice sets, generating

[1] We could have defined choice sets as segments defined by an arbitrary set of job characteristics, as Şahin, Song, Topa, and Violante (2014) and Herz and Van Rens (2020), or by clustering algorithms, as Banfi and Villena-Roldán (2019). The notion of "local labor market" in Labor Economics seems largely attached to geography (Moretti, 2011; Foote, Kutzbach, and Vilhuber, 2021), which seems insufficient for dealing with objects of high-dimensional attributes such as job ads.

key variability for parameter identification.

We then run a linear probability model to explain whether a job seeker applies or not to an ad in her consideration set. On the right-hand side, we put demographic covariates and different gaps between applicants' characteristics and job ad requirements in continuous variables such as years of education, experience, log wages, and distance, as well as binary characteristics such as occupation. Since the composition of applicants and ads may have substantially changed over time due to the penetration of the website and cyclical factors, we use the DiNardo, Fortin, and Lemieux (1996) reweighting technique to keep the composition fixed at 2016 Q3, the period with the largest participation in our database. Our weighting method works as long as applicants are unaware of the specific identities or characteristics of other applicants on online job boards, a realistic description of this market. Using a specific applicant as an example, the computed weights for her applications involve the choices of other applicants who ignore the existence of the reference applicant.

With respect to the *intensive margin*, our empirical exercise reveals the effects of different traits on the probability of applying to a job ad. Males apply more than females, especially if they are unemployed. Marital status matters: employed married individuals apply less than their single counterparts. We also document that the application probability increases with age for the unemployed but shows a mild decay for the employed. Although the job finding probability declines with age (Choi, Janiak, and Villena-Roldán, 2015; Menzio, Telyukova, and Visschers, 2016), both findings may be consistent as long as older jobseekers react to scarcer job opportunities and a higher opportunity cost of time.

We find some evidence of stock-flow matching behavior, in the vein of Taylor (1995); Coles and Muthoo (1998); Coles and Smith (1998) and Ebrahimy and Shimer (2010), among others: new job seekers on the website (the flow) apply to the stock of job ads during their initial time on the platform. When time passes, the inflow of job seekers becomes part of the stock of individuals, who then try to match the new flow of job positions. Gregg and Petrongolo (2005) and Coles and Petrongolo (2008) provide evidence for stock-flow matching in different contexts. In our analysis, we show that unemployed individuals apply more for newer job ads, other things being equal, which is consistent with their distaste for "phantom" job postings (Chéron and Decreuse, 2016; Albrecht, Decreuse, and

Vroman, 2023). However, we observe no significant effect for on-the-job seekers. On top of this, we show that workers apply more to job postings advertising more vacancies, suggesting that individuals respond to indications of the likelihood of receiving an offer given an application.

We also find a decreasing probability of application as the unemployed search duration increases, an important issue for the design of unemployment insurance policies, as stated by Faberman and Kudlyak (2019) and DellaVigna, Heining, Schmieder, and Trenkle (2021), among others. For on-the-job seekers, we find that the average application probability is mildly decreasing and then relatively flat with respect to tenure in the current job, which is somewhat consistent with theory (Jovanovic, 1979; Li and Weng, 2017) and previous evidence (Pissarides and Wadsworth, 1994; Fujita, 2012). This also seems qualitatively consistent with the evidence of realized job-to-job flows in Menzio, Telyukova, and Visschers (2016). This finding is relevant to discipline models explaining job-to-job transitions and frictional wage dispersion, as in Hornstein, Krusell, and Violante (2011).

We measure the cyclical behavior of the intensive margin with respect to regional unemployment rates of the applicant's residence because our data does not cover a long period of time, an approach also followed by Hazell and Taska (2020). Our finding is that for relatively low regional unemployment rates, below 6%, the application probability does not change much; it shows strong procyclicality beyond that level. The behavior of employed and unemployed job seekers seems very similar in this regard, suggesting that the search effort cannot offset the relative scarcity of job opportunities in slack regional labor markets. Our findings may help reconcile disparate results in this literature that show that effort is either procyclical (Gomme and Lkhagvasuren, 2015), acyclical (Leyva, 2018), or even countercyclical (Mukoyama, Patterson, and Şahin, 2018).

As for the *selective margin*, we find that applicants show *misalignment* distaste in general: they are highly sensitive to the fit between worker traits and ad requirements in terms of educational level, experience, location, wages, and occupation. This distaste is apparent in terms of location, i.e., the probability of application decreases with the geographical distance between worker and employer, and in terms of occupation. The results are nuanced in other dimensions: job seekers target their most preferred type of job on average, but it is not necessarily the one that matches their current characteristics. For instance, all workers, especially the employed, possess substantially more experience than the minimum required by the applied ad, on average.

In the evidence we interpret that employed job seekers are more daring or *ambitious* than the unemployed. On-the-job seekers apply more frequently to jobs for which they are slightly underqualified in terms of education and to jobs with wages above their expectations on average. In contrast, the unemployed apply the most to jobs that have the educational requirements more closely aligned with those they possess and to ads that match more closely their own wage expectations. Thus, the evidence may be consistent with unemployed seekers maximizing job offer chances while employed ones are trying to climb the job ladder.

Our paper is generally related to a growing literature that uses data from online job posting and search websites in order to study different aspects of frictional markets. Kudlyak, Lkhagvasuren, and Sysuyev (2013) study how job seekers direct their applications over the span of a job search. They find some evidence on the positive sorting of job seekers to job postings based on education and how this sorting worsens the longer the job seeker spends looking for a job (the individual starts applying for worse matches). Faberman and Kudlyak (2019) use online job board data to study the intensive margin of job search. Marinescu and Rathelot (2018) use information from `www.careerbuilder.com` and find that job seekers are less likely to apply to jobs that are farther away geographically. Marinescu and Wolthoff (2020) use the same job posting website to study the relationship between job titles and wages posted on job advertisements. They show that job titles explain nearly 90% of the variance in explicit wages. Banfi and Villena-Roldán (2019) and Banfi, Choi, and Villena-Roldán (2022) use data from this website to find substantial evidence of directed search and assortative matching, providing complementary evidence related to the selective margin. Fluchtmann, Glenny, Harmon, and Maibom (2023) merge administrative data and online job board applications to study the dynamics of applied-for wages for the unemployed rather than search intensity as we do.

2 The data

We use data from `www.trabajando.com` (henceforth the website), a job search engine operating in Chile, covering a sample of job postings and job seekers between January 1st 2008 and December

24th, 2016. The raw information in the dataset contains more than 14 million single applications, from around 1.5 million job seekers to around 270 thousand job ads.

Our dataset has detailed information on both applicants and recruiters. First, we observe entire histories of applications from job seekers and dates of ad postings (and repostings) for recruiters. Second, we have detailed information for both sides of the market. For job seekers, we observe date of birth, gender, nationality, place of residency ("comuna" and "región", akin to county and US state, respectively), marital status, years of experience, years of education, college major, and name of the granting institution of the major.[2] We have codes for the occupational area of the current or last job of individuals,[3] information on their salary and both their starting and ending dates.

In terms of the website's platform, job seekers can use the site for free, while firms are charged for posting ads. Job advertisements are posted for a minimum of 60 days, but firms can pay additional fees to extend this term.[4]

For each posting, we observe its required level of experience (in years), required college major (if applicable), indicators on required skills (specific, computing knowledge and/or "other"), how many positions must be filled, the same occupational code applied to workers, geographic information ("región" only) and some limited information on the firm offering the job: its size (number of employees in brackets) and industry (1 digit code).[5] Educational categories are *primary* (one to eight years of schooling), *high school* (completed high school diploma, 12 years), *technical tertiary education* (professional training after high school, usually 2-4 years), *college* (completed university degree, usually 5-6 years) and *post-graduate* (any schooling higher than a college degree).

A novel feature of the dataset, compared to the rest of the literature, is that the website asks job seekers to record their expected salary, which they can then choose to show or hide from prospective employers. Recruiters are also asked to record the expected pay for the job posting and are given the same choice as to whether to make this information visible to the applicants. Naturally, one could

[2]This information is for any individual with some post-high school education.

[3]We observe a one-digit classification, created by the website administrators. These codes are not comparable to standardized classification taxonomies such as SOC or ISCO.

[4]As of January 3rd, 2018, the 60-day fee is CLP 69,900 + 19% VAT, which is equivalent to USD 136 or EUR 113. There are quantity discounts for big clients, too.

[5]We observe an industry classification created by the website administrators that does not match formal taxonomies such as NAICS or ISIC.

question the reliability of wage information, which will ultimately be hidden from the other side of the market. Banfi and Villena-Roldán (2019) address the potential issue of "nonsensical" wage information in job ads by comparing the sample of explicit vs. implicit (job ads without any salary information) postings by firms and find that observable characteristics predict fairly well implicit wages and vice versa. Moreover, even if employers choose to hide wage offers, they are used in filters of the website for applicant search. Hence, employers are likely to report accurately even if their wage offers are not shown because misreporting may generate adverse consequences. On the other hand, a major caveat of our dataset is the absence of information on activities performed outside the website, such as individuals seeking jobs through other means and, more importantly, the outcomes of job applications.

For the remainder of the paper, we restrict our sample to only individuals working under full-time contracts and those who are unemployed. We further restrict our sample to individuals aged 23 to 60. We discard individuals reporting desired net wages above 5 million pesos.[6] This amounts to approximately 8,347 USD per month[7], which is higher than the 99th percentile of the Chilean wage distribution, according to the 2013 CASEN survey.[8] We also discard individuals who desire net wages below 159 thousand pesos (around 350 USD) a month (the legal minimum wage at the start of our considered sample). Consequently, we also restrict job postings to those offering monthly salaries within those bounds.

Our units of analysis are individual *applications*. We restrict our sample to individuals who were actively looking for a job (i.e., made an application) and job postings that received at least one application. While we observe long histories of job search for a significant fraction of workers (some workers have used the website for several years), we consider only applications pertaining to their last job search "spell", which we define as the time window between the last modification/creation of their online curriculum vitae (cv) on the website and the time of their last submitted application

[6]In the Chilean labor market, wages are usually expressed in a monthly rate net of taxes and mandatory contributions to health (7% of monthly wage), to fully funded private pension system (10%), disability insurance (1.2%), and mandatory contributions to unemployment accounts (0.6%)

[7]Using the average nominal exchange rate between 2013-16 at the Central Bank of Chile statistics website.

[8]CASEN stands for "Caracterización Socio Económica" (Social and Economic Characterization), and aims to capture a representative picture of Chilean households. For data and information in Spanish, visit this website.

or the one year mark, whichever happens first. Since individuals maintain information about their last job in their online profile, as well as contact information and salary expectations, we assume that any modification of this information is done primarily when individuals who are currently working or who have already used the website in the past are ready to search in the labor market again. We cannot infer any labor transition based on application behavior because employed individuals may keep searching for jobs, and unemployed individuals may search outside of the website. We further drop individuals who apply to more than the 99-th percentile of job applicants in terms of number of submitted applications in the defined window.

Table 1 shows descriptive statistics for the job seekers in our sample. From the table, we observe that the average age is 33.5 and that job seekers are comprised of mostly single males, with 59.71% being unemployed ($128,482$ unemployed seekers from a total of $215,169$ individuals.). Average experience hovers around eight years. Job seekers in our sample are more educated than the average in Chile, with 41.84% of them having a college degree, compared to 25% for the rest of the country in the comparable age group (30 to 44 years of age), according to the 2013 CASEN survey. There is also a big discrepancy by labor force status: unemployed seekers are significantly less educated on the website.

From the table, we can also observe that most job seekers claim occupations related to management (around 20%) and technology (around 25%) and that average expected wages are approximately (in thousands) CLP\$ $1,087$ and CLP\$ 592 for employed and unemployed seekers, respectively. For comparison, the 2013-16 average minimum monthly salary in Chile was around CLP \$ 226 thousand.[9]

In terms of search activity, the average search spell amounts to around five weeks. The amount of time spent searching for a job is higher for those employed than for the unemployed: 5.24 versus 4.83 per week, respectively. In terms of applications, both groups show very similar choices, with around 1.52 submitted applications. The unemployment spell linked to an application refers to the elapsed time between the separation date of the last job and the actual application date.

[9]The minimum wage has increased substantially in recent years. For information about the trajectory of the legal minimum wage in Chile, please see this website.

Table 1: Characteristics of Job Seekers

	Employed	Unemployed	Total
Demographics (%)			
Male	62.03	53.97	57.21
Married	33.80	27.50	30.03
Demographics (Avg)			
Age	33.77	33.25	33.46
Experience (years)	8.28	7.64	7.90
Wages (thousand CLP)	1,087	592	792
Tenure (weeks)	179.29	–	179.29
Unemployment duration (weeks)	–	60.17	60.17
Education level (%)			
Primary (1-8 years)	0.12	0.25	0.2
High School	17.94	36.89	29.25
Technical Tertiary	26.56	28.82	27.91
College	54.22	33.48	41.84
Post-graduate	1.17	0.55	0.8
Occupation (%)			
Management	23.5	17.85	20.12
Technology	31.59	21.21	25.39
Not declared	20.29	42.54	33.57
Rest	24.62	18.4	20.92
Search Activity			
weeks searching on website	5.24	4.81	4.98
Number of applications	1.47	1.51	1.50
Observations	86,687	128,482	215,169

3 Application probabilities and job seeker preferences

In this section, we analyze empirically which attributes of heterogeneous jobs attract more applications from heterogeneous job seekers. To do this, we first need to determine which set of job ads is relevant for each individual in our sample. However, our dataset only contains information on actual applications, and no information is collected by the website on the total number of searches or *clicks* on job postings by individuals. Thus, we do not have sample variation in terms of job ads: we only observe those that individuals choose to apply to, not those that are observed but then discarded by seekers. This problem of "consideration sets" (i.e., the set of products consumers are aware of) is addressed in the literature in marketing and industrial organization (Van Nierop, Bronnenberg, Paap, Wedel, and Franses, 2010; Abaluck and Adams-Prassl, 2021), but our approach is essentially different. Nevertheless, we hypothesize that our network-revealed preference approach could be used in this literature, provided databases identify the goods purchased by each consumer.

3.1 Market segmentation through network analysis.

We could consider the cross between all job seekers and all job ads that are time feasible in our sample, what we call the *exploded* dataset. However, there are major drawbacks to this approach: First, the exploded dataset makes comparisons between job seekers and job positions, which may be objectively too different to consider. A typical job seeker may find more than 20,000 available job ads for her to screen and choose from, implying an unrealistic effort for workers. Second, since we truly try to characterize an actual decision-making process, by including job ads never considered by the applicant in her choice set, we introduce an attenuation bias because most of the included job ads actually have zero probability of being applied to. Third, a more practical issue is that the size of the estimating sample becomes simply too large to handle,[10] making the task of even simple calculations infeasible.

We could create choice sets by clustering job ads using their traits, as in Banfi and Villena-Roldán (2019). However, such an approach links workers to a fixed set of ads with little or no cross-sectional variation across similar applicants. Instead, our approach uses the revealed preferences of workers to construct individual consideration or choice sets based on coincidental choices made by other applicants. In reality, workers can potentially apply to jobs considering a large number of potential characteristics, many of which we can observe. The revealed preference approach circumvents the problem of defining what the relevant information is for workers. We only determine the relevant set of ads for each worker based on their actual choices, regardless of the way they process their information.

To formalize this notion, we use the network formed by job seekers to determine which job postings are relevant to them. Assume that each individual represents a node in the network and that a link between nodes is defined as *having applied to the same job posting*. For each job seeker w, we

[10]With our sample constraints, we have 215,000 workers who could potentially apply to 20,000 job ads when they change their CV, a very conservative lower bound if they stay actively applying for several weeks. Thus, the exploded set contains $215,000 \times 20,000 = 4,3$ billions of potential applications.

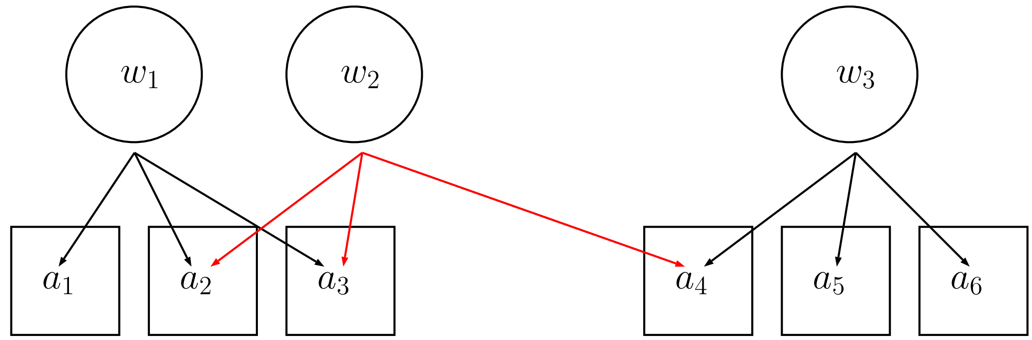

Figure 1: Example of a network formed by workers $\{w_1, w_2, w_3\}$. Worker w_1 is linked to worker w_2 by common applications to ads a_2 and a_3 but is not linked with w_3 in the network of degree 1. All workers are linked in the network of degree 2.

define the set of relevant job postings $\mathcal{A}_w^1$ as the union of all job postings applied by the set of all job seekers linked to w. This is what we define as a network of degree 1, since for each individual, we only consider their immediate links (1 degree of separation).

Following this logic, the network of degree 0 is the original dataset for individual w ($\mathcal{A}_w^0$), since the network contains only information about job seekers and their applications (no information on links is used). On the other hand, a network of degree 2 is defined as the network that considers both job seekers linked directly to w in addition to those who are linked with the links of w (job seekers have 2 degrees of separation), giving rise to the set $\mathcal{A}_w^2$. We can continue with this logic iteratively until we form the set $\mathcal{A}_w^\infty$, which is the cross between each job seeker w and all job postings a as long as they are connected somehow through the network.[11]

Figure 1 shows an example of the network algorithm and the resulting datasets. In the figure, there are three workers, $\{w_1, w_2, w_3\}$ and six job postings, $\{a_1, a_2, a_3, a_4, a_5, a_6\}$. Consider worker w_1. She has applied to three jobs, thus $\mathcal{A}_{w_1}^0 = \{a_1, a_2, a_3\}$ and is linked to w_2 through applications to $\{a_2, a_3\}$. Since w_2 also applied to job position a_4, one can infer that some characteristics of a_4 are not desirable to w_1. If we consider networks of degree 1, a_4 would be included in the set of relevant ads for the first worker. Notice also that in this example, w_1 is not directly linked with w_3, or in our language, the degree of separation between these two workers is higher than 1.

Again, considering the first worker, we have $\mathcal{A}_{w_1}^0 = \{w_1, w_2, w_3\}$, and as discussed above, $\mathcal{A}_{w_1}^1 =$

[11] Technically, the set $\mathcal{A}_w^\infty$ and the exploded dataset differ if there are isolated pairs or groups of individuals who are not connected to the rest of the applicants through any ad.

$\{a_1, a_2, a_3, a_4\}$. Given that w_1 and w_2 are linked and that w_2 is linked with w_3, the relevant job ads for w_1, given a network of degree 2, is $\mathcal{A}^2_{w_1} = \{a_1, a_2, a_3, a_4, a_5, a_6\}$. In our simple example, the network of degree 2 is already the "exploded" network (all ads to all workers).

The formal definition of a one-degree-of-separation ad set for a worker w is

$$\mathcal{A}^1_w = \bigcup_{v:\mathcal{A}^0_w \cap \mathcal{A}^0_v \neq \emptyset} (\mathcal{A}^0_w \cup \mathcal{A}^0_v)$$

which can be generalized for other degrees of separation.[12] In what follows, we will concentrate on networks of degree one only.

Table 2: Number of relevant ads (a) per worker (w)

	Potential ads for a worker		
	All	U	E
percentile 10	2	2	2
percentile 50	16	16	19
percentile 90	96	104	87
mean	38.5	40.7	36.8
standard deviation	68.1	73.8	57.1
mean applications (%)	22.3	23.2	20.9

Notes: The table shows the number of relevant job postings per job seeker given a network of degree 1 (see main text). Statistics separated by labor force status of job seeker (U = unemployed, E = employed). Standard errors in parentheses. One, two, and three asterisks indicate significance at 10%,5%, and 1%, respectively.

In table 2, we present information on the resulting number of relevant job postings per worker and workers per job posting, given a network of degrees one. The median number of relevant job postings (a) is 16 per job seeker, with employed seekers being related to more posts (19) than those unemployed (16). The number of potential ads exhibits quite a bit of variation, going from 2 (tenth percentile of distribution) to 104 and 89 for the unemployed and employed, respectively (ninetieth

[12] The generalization follows a recursive definition

$$\mathcal{A}^s_w = \bigcup_{v:\mathcal{A}^{s-1}_w \cap \mathcal{A}^0_v \neq \emptyset} (\mathcal{A}^{s-1}_w \cup \mathcal{A}^0_v)$$

which depends on $\mathcal{A}^0_w$ and the definition of $\mathcal{A}^1_w$.

percentile). Given the sets of related job ads, mean application rates,[13] are 22.3% for the entire sample, with unemployed seekers applying to 23.2%, while employed ones do so for 20.9% of their relevant ads.

Although our network approach allows us to build choice sets for applicants, it is unlikely that all the ads in the set are equally considered.[14] In any given network-induced set of choices for each worker w, there is heterogeneity in the relevance of job ads according to how strong the link between two workers is. Intuitively, the bigger the overlap in submission choices by both workers, the closer they are and the more relevant the additional job ads are for each other. As an example, consider worker w_2 in figure 1. Since w_2 and w_1 submit common applications to several common positions, they likely have similar preferences and qualifications. Then, the likelihood that w_2 truly considers applying to the job ads to which w_1 applied must be high. In contrast, worker w_2 and worker w_3 share fewer applications, so the likelihood that w_2 considered $\{w_5, w_6\}$ is lower.

To give more formality to this intuition, we construct a weight function $q(a, w)$ for each worker w and job position a. In order to construct q, we start with the function $b(w, v)$ which we apply to all pairs of linked workers w and v, as a measure of how similar they are in terms of application decisions. We construct b, given some general restrictions:

1. $b(w, v) \in [0, 1]$
2. $b(w, w) = 1$
3. $b(w, v) = 0$ if and only if $\mathcal{A}_w^0 \cap \mathcal{A}_v^0 = \emptyset$

On top of conditions 1–3 above, we want the function $b(w, v)$ to be monotonic in set similarity. A particular functional form that satisfies these conditions is

$$b(w, v) \equiv \frac{|\mathcal{A}_w^0 \cap \mathcal{A}_v^0|}{|\mathcal{A}_w^0 \cup \mathcal{A}_v^0|} \tag{1}$$

[13] Defined as the number of effective applications to total ads for worker w:

$$\frac{|\mathcal{A}_w^0|}{|\mathcal{A}_w^1|}$$

[14] Moreover, some ads can be unavailable for some workers due to different time frames for search windows. We discard these ads when appropriate.

where $|S|$ is the cardinality (number of elements) of set S. Equation (1) is also known as the *Jaccard Similarity Index* between two groups (Jaccard, 1901).

We then define the weight of an ad a for a worker w as:

$$q(a,w) = \max_{v:\, a \in\, \mathcal{A}_v^0} \{b(w,v)\} \tag{2}$$

Intuitively, we consider the importance of a particular job ad a in the choice set of w given the similarity between the choice set of w and the most similar choice set of any other applicant v, linked to w. It is easily verified that, given this proposed weighting function:

1. $q(a,w) = 1$ if and only if $a \in \mathcal{A}_w^0$. (w applies to a)
2. $q(a,w) \in [0,1)$ if and only if $a \notin \mathcal{A}_w^0$ (w does not apply to a)
3. $q(a,w) = 0$ if and only if $a \notin \mathcal{A}_w^1$ (a is not in the choice set of w)

These definitions operationalize a similarity notion between workers and ads that we use to estimate application equations. Essentially, we construct $q(a,w)$ using the choices of applicants who are not w and whose characteristics or applications are unknown to w.

3.2 Estimation of the application equation

For the constructed dataset, we estimate preferences of job seekers based on their observed characteristics and the ones posted by ads that are relevant to them. More specifically, we estimate a linear regression of the form

$$y_{aw} = X_{aw}\beta + \sum_k \sum_{p=1}^{\bar{P}} \{\gamma_{kp}(z_{k,aw})^p\} + \gamma_j z_{j,aw} + \sum_k \sum_\ell \mathbf{1}_{\{k\neq\ell\}} \delta_{k\ell} z_{k,aw} z_{\ell,aw} + \epsilon_{aw} \tag{3}$$

where y_{aw} is a dummy variable that takes the value of one if a job seeker w applies to posting a and zero otherwise. In X_{aw}, we include a linear trend and monthly dummies to control for secular trends and seasonal patterns in website usage and penetration. We also control for observed job and worker characteristics, which do not overlap. The list of variables for the job includes firm size, dummies for firm industry, specific job requirements (computer knowledge or some other form of specific knowledge), and controls for specific job characteristics: type of contract (full or part time), number of vacancies needed to be filled, and controls for job title relevant words,[15] following Marinescu and Wolthoff (2020) and Banfi and Villena-Roldán (2019). For individuals, we control for marital status (dummy variable for marriage), gender (dummy for male), an interaction between married and males and quintic polynomials for the age of the job seeker, and the amount of time (measured in weeks) in either the current job, a tenure (for those employed), or unemployment (for unemployed seekers). Among the employed, 40.4% of the sample have no measured tenure since the starting date of the job is unreported. To keep these observations in our sample, we define a dummy variable for missing tenures and impute a value of zero to all unobserved tenures. In this way, the estimated tenure profile should be interpreted as conditional on declaring a starting date for the current job. The missing variable coefficient, in turn, is the differential effect in application probability of an undeclared starting date with respect to an observed zero tenure. The same strategy is used for unemployed job seekers, but in this case only 7.7% of starting dates are missing.

For both seekers and ads, we include a variable indicating whether the wage expectation (for seekers) or the wage expected to be paid (for jobs) is made explicit or not. To control for business cycle conditions, we consider the unemployment rates of the applicant's region during the month in which the application took place.[16] We also include a quadratic term of the regional unemployment rate to capture potential non-linear effects. Since our sample does not cover an extended period, using region cross-regional variation to identify the cyclical effects is a sensible approach, as do Hazell and Taska (2020). The effects of these characteristics impact the level of the probability of application and therefore are related to the *intensive margin* of the application process that has been more profusely studied in the literature (DeLoach and Kurt, 2013; Gomme and Lkhagvasuren, 2015; Baker and Fradkin, 2017; Ahn and Shao, 2017; Leyva, 2018; Mukoyama, Patterson, and Şahin, 2018; Faberman and Kudlyak, 2019; Bransch, 2021).

[15] As customary in text analysis, we drop stop words that are either too frequent or convey no specific meaning.

[16] For worker-ad pairs that are matched given our network algorithm, the date of an actual application does *not* exist. In those cases, we impute the date of application by the mode date of applications of the linked workers to that particular job ad.

On the other hand, to provide a novel measurement of the *selective margin* of job search, we include a set of controls for the *misalignment* or gap (which we denote by z) between characteristics required by firms and the characteristics of the job seeker. For continuous variables, which we index by k, we define z_k as the simple difference between the value of the characteristic required by the position and the value of the characteristic possessed by the job seeker. We do this for years of education, years of experience, and log wages. For regional distance, we compute misalignment as hundreds of kilometers between regional capital cities, using applications to job ads that have only one reported region, so we do not consider ads offering jobs with unknown, multiple, or international locations. We observe that 30.2% and 27.7% of applications done by employed and unemployed workers, respectively, involve a job ad with no unique Chilean region. To avoid losing these observations, we define a dummy variable for missing ad region, and impute a value of zero to all unobserved regional distances. Hence, the estimated distance application profile is conditional on the job ad declaring a region. The missing ad region coefficient captures the differential effect in the application probability of a missing ad region with respect to an observed zero distance.

For occupations, the variable z_j is defined as a dummy that takes the value of 1 when the category in the job posting is different from the characteristic of the worker and 0 when they are the same.

In equation (3), for each of the continuous dimensions k we include in the regression a polynomial of order $P = 5$ to assess whether non-linearities exist in the effect of these *misalignments* on application decisions. In this way, we capture if *over*-qualified ($z_k < 0$) jobseekers behave differently from *under*-qualified ($z_k > 0$) ones. We estimate the above equation, separating our sample between the employed and unemployed, to assess whether on-the-job search behavior differs from unemployed search behavior. We also consider interaction effects between different *misalignment* levels.

On top of weighting observations to reflect market segmentation, we need to take into account changes in the market composition of both applicants and job ads. Since we drop all applications made before the last CV update, our sample disproportionably covers individuals from years close to the temporal end of our sample. Indeed, 24.5% of our data, or about 2 million observations, are applications made in the third quarter of 2016. Balancing the composition is important if we want to disentangle cyclical search behavior from compositional effects in the data. This is quite important due to the website's increasing penetration in the Chilean labor market. We address this issue by using the reweighing technique of DiNardo, Fortin, and Lemieux (1996). We choose the composition of jobs and workers in 2016 Q3 and run a probit model estimating the probability of an application occurring in 2016 Q3 as a function of observables on the applicant side and on the job ad side[17]. We compute predicted probabilities $\widehat{p}(a,w)$. In our results, we define a final weight for worker-ad observations as

$$\varphi(a,w) = q(a,w)\frac{\widehat{p}(a,w)}{1-\widehat{p}(a,w)}$$

for applications in the common support of observations from 2016 Q3 and those from other quarters.[18] In the latter formula, we remind that $q(a,w)$ is the market segmentation weight we described in section 3.1.

Technically, for a consistent estimation, we require that the network part of the weight φ, $q(a,w)$ and the error term ϵ in equation (3) be zero conditional on the observables of the pair (a,w), i.e., $cov(\sqrt{q},\epsilon|X,\{z_{k,aw}\}_k,z_j) = 0$.[19] As shown in section 3.1, q essentially reflects choices made by anonymous co-applicants of w, who generally ignore worker w existence, not to mention her unobserved characteristics. This exactly describes the operation of online job boards: typically, no one knows who else is applying for the same job one is applying to. However, a caveat is that the residual term ϵ may be correlated with ad unobserved features. Since we control for a very large array of ad characteristics, this may only occur if the applicant w and some others possess private information about the posted job that is not disclosed in the ad.

[17]For categorical dummies, we drop those whose average for 2016 Q3 or for the other quarters was below 0.2 or above 0.8 for some group to avoid very high or very low predicted probabilities

[18]For the sake of completeness, we report the results without using DFL weights in the appendix.

[19]A more detailed explanation of this result is as follows. Consider that all the right-hand side variables comprising worker w and ad a characteristics $X_{aw},\{z_{k,aw}\}_k,z_{j,aw}$ are stacked into the vector S_{aw} and all the corresponding parameters are stacked into the vector θ so that the linear probability model would be $y_{aw} = S_{aw}\theta + \epsilon_{aw}$. Then, using weights $q(a,w)$, we run regressions of the form $\sqrt{\varphi(a,w)}y_{aw} = \sqrt{\varphi(a,w)}S_{aw}\theta + \sqrt{\varphi(a,w)}\epsilon_{aw}$ to estimate coefficients weighted by q, as in a standard version of Generalized Least Squares. To ensure that we obtain consistent estimators, it must be true that $E[\sqrt{\varphi(a,w)}\epsilon_{aw}|S_{aw}] = \sqrt{\frac{\widehat{p}(a,w)}{1-\widehat{p}(a,w)}}cov(\sqrt{q(a,w)},\epsilon_{aw}|S_{aw}) = 0$, which occurs only occurs if the last covariance equals zero. A stronger condition is to require independence between weights $q(a,w)$ and the application error term ϵ.

4 Results

Table 3 shows coefficients multiplied by 100 from the estimating equation (3) using ordinary least squares. We report estimates by employment status and whether we perform network and DFL weighting or not of our estimates. Unweighed estimates typically do not change signs, but their magnitudes are attenuated. A possible interpretation is that our weighing method reduces the importance of irrelevant job ads into applicants' consideration sets.

4.1 Results on the intensive margin: Applicants and ad traits

Table 3: Intensive margin coefficients by labor status

VARIABLES	(1) Employed weights	(2) Employed no weights	(3) Unemployed weights	(4) Unemployed no weights
Married	-2.646***	-0.488	0.373	0.004
	(0.928)	(0.366)	(0.495)	(0.203)
Male	1.583***	0.262***	2.301***	0.352***
	(0.065)	(0.031)	(0.050)	(0.025)
Explicit wage (w)	0.224***	-0.131***	0.748***	-0.062***
	(0.057)	(0.028)	(0.047)	(0.024)
Explicit wage (a)	-2.767***	-0.793***	-0.965***	-0.106***
	(0.088)	(0.042)	(0.067)	(0.032)
No. of Vacancies (a)	-0.011	-0.003	0.039***	0.006***
	(0.011)	(0.002)	(0.006)	(0.001)
Ad duration (weeks)	-0.002	0.003**	-0.085***	-0.020***
	(0.002)	(0.001)	(0.001)	(0.001)
Observations	2,124,294	2,124,294	3,184,960	3,184,960
R-squared	0.119	0.037	0.110	0.038
Mean app prob	22.59	4.391	23.17	4.698

Notes: Regression coefficients from a linear regression on application decisions. Dependent variable is y_{aw}, a dummy for the existence of a job application. Each regression controls also for polynomials and interactions in *misalignment* as well as age of the worker, firm size, contract type, dummies for different types of requirements of the job and characteristics of the firm (see details in the main text). Standard errors in parentheses. One, two, and three asterisks indicate significance at 10%,5%, and 1%, respectively.

The evidence shows that the unemployed job seekers apply more frequently then the employed to the ads in their consideration set. Among the employed married individuals apply less than non-married counterparts, while there is a non significant gap for the unemployed. Male job seekers, especially unemployed ones, apply more often to ads into their consideration sets, keeping other applicant and ad characteristics constant.

Individuals who choose to be explicit about their wage expectations at the time of an application apply more often than those hiding wage expectations, especially if unemployed. The table also shows that an explicit wage in the job ad negatively affects the decision to apply on average, but the effect is weaker for the unemployed. This is consistent with findings in Banfi and Villena-Roldán (2019) who show that ads with hidden wages tend to attract more applicants due to a higher likelihood of potential wage flexibility or bargaining, as suggested in the Michelacci and Suarez (2006) model. Moreover, unemployed individuals significantly increase their probability of application by 0.04% when the n number of vacancies increase in one, while the employed has no significant reaction. The small positive response to a marginally higher likelihood of receiving an offer suggests an important role for selection on the employer side, i.e. non-sequential employer search (van Ours and Ridder, 1992; van Ommeren and Russo, 2013) or signaling for poorer job conditions.

The effect of the perceived "age" of the job ad has a negative effect for the unemployed, whodislike job ads that are older (in weeks). The negative effect for the unemployed can be related to stock-flow matching behavior[20]: new job seekers in the website (the flow) apply to the stock of job ads. When time passes, the inflow of job seekers becomes part of the stock of individuals, who then try to match with the new flow of job positions, as suggested by evidence in Gregg and Petrongolo (2005) and Coles and Petrongolo (2008). Our results for the unemployed are also consistent with applicants reacting to "phantom" ads, which may be filled positions by the time of the potential application, as in Albrecht, Decreuse, and Vroman (2023) and Chéron and Decreuse (2016). The evidence reported by Davis and Samaniego de la Parra (2017) is also qualitatively consistent with our findings. The effects are not significant for the employed, which suggests a different pattern of search for this group in this website.

[20]References are Taylor (1995); Coles and Muthoo (1998); Coles and Smith (1998); Ebrahimy and Shimer (2010)

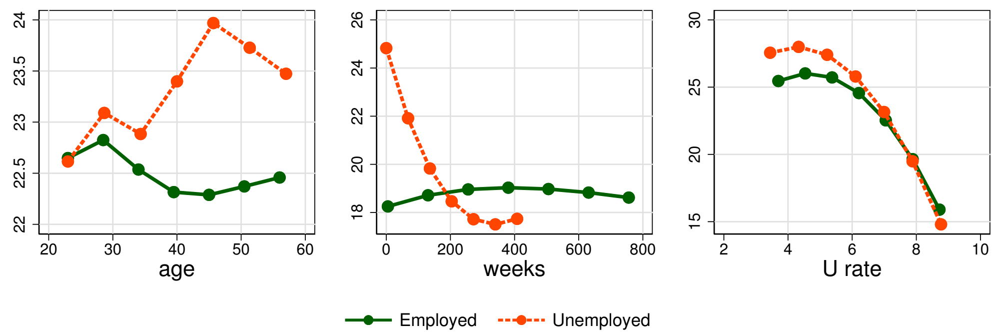

Figure 2: Predicted application probabilities for different ages, number of weeks in the current labor force status, and national unemployment rate at the time of the application decision, given results from equation (3). The figure is computed using the coefficients associated to a polynomial of order 5 on each variable and leaving the rest of regressors at their sample mean.

4.2 Results on the intensive margin: Life-cycle, duration and business cycle effects

We report the predicted application probability varying age, duration of employment status, and unemployment between the 5th and 95th percentiles of their sample values, while keeping the other covariates at their mean values in figure 2.

In the left panel, we observe that the unemployed apply more often to ads in their consideration sets at all ages, and their probability of application increases with age, with an overall peak at age 45 to decrease until mid-fifties. For the unemployed, the application probability is higher for individuals under 30, and then decreases until the mid-forties, and then slightly increases. While this evidence might seem only partially consistent with job finding rates and employment-to-employment transitions over the life-cycle as reported by Choi, Janiak, and Villena-Roldán (2015) and Menzio, Telyukova, and Visschers (2016), and Naudon and Pérez (2018) for Chile, we point out two reasons why this is not the case. First, in these papers, job finding rates refer to the larger frequency of realized transitions. In contrast, our evidence here is about search or application effort. Indeed, Mukoyama, Patterson, and Şahin (2018) show a slightly increasing profile of effort on the intensive margin of time devoted to job search until age 50. Second, the self-selected sample of older workers using the online job board may be somewhat different from the average worker in the labor force of that age.

The middle panel in the figure shows a decreasing application probability as the search duration increases, measured as the time elapsed between the finishing date of the previous job and the application date. The extended range of durations suggests that equalizing traditional unemployment duration with our measure of search duration is far-fetched. Thus, an appropriate interpretation is that individuals who have lost jobs and are website users make most of their applications soon after the separation. For employed jobseekers, the application likelihood seems to be flat for the most part even though there is a slightly increasing trend up to 400 weeks, or nearly eight years. Two offsetting factors may be at play: a growing match-specific human capital deterring on-the-job search and a market-learning process of the worker that increases the outside value of the applicant.

In terms of business cycle conditions, the right panel of figure 2 shows a clearly decreasing relationship between the unemployment rate, our cyclical variable, and application decisions. The application probability remains quite flat between 25-30% and slightly higher for the unemployed when the regional labor market exhibits low unemployment rate, i.e. below 5.5%. When moving to regional labor markets showing unemployment rates between 6% and 8.3%, the average application probability declines from 25% to 15% and is very similar for both employed and unemployed applicants. The decreasing pattern for above-the-median unemployment rate suggests that job seekers find that their search effort cannot compensate for the scarcity of available jobs when unemployment is high, unlike Faberman and Kudlyak (2019), Mukoyama, Patterson, and Şahin (2018) and Bransch (2021). In contrast, the finding aligns with DeLoach and Kurt (2013) and Gomme and Lkhagvasuren (2015). Yet Leyva (2018) finds roughly acyclical search effort, as in the lower-end of the unemployment rate in our sample. Our finding of non-monotonicity of the effect helps reconciling these heterogeneous pieces of evidence in the literature.

4.3 Selective Margin: Misalignment and applications.

We present the effect of *misalignment* in continuous dimensions (education, experience, log wages, and distance). As noted above, for each dimension, we take the simple difference between what is required in the job ad (years of experience, for example) and what the job seeker possesses. Then, a negative value for this misalignment measure means that the individual is "overqualified" in the particular dimension and that the individual is "underqualified" if it is positive. For instance, if an applicant has more years of experience than the minimum required in the job ad, she is "overqualified" in terms of experience. With some abuse of language, we will define that a worker applying to an ad posting a wage lower than her own expectations is "overqualified". For geographical distance, these notions are not applicable, so we concentrate only on the absolute distance (in hundreds of kilometers) between the capital of the region of the job and the capital of the region of the applicant.

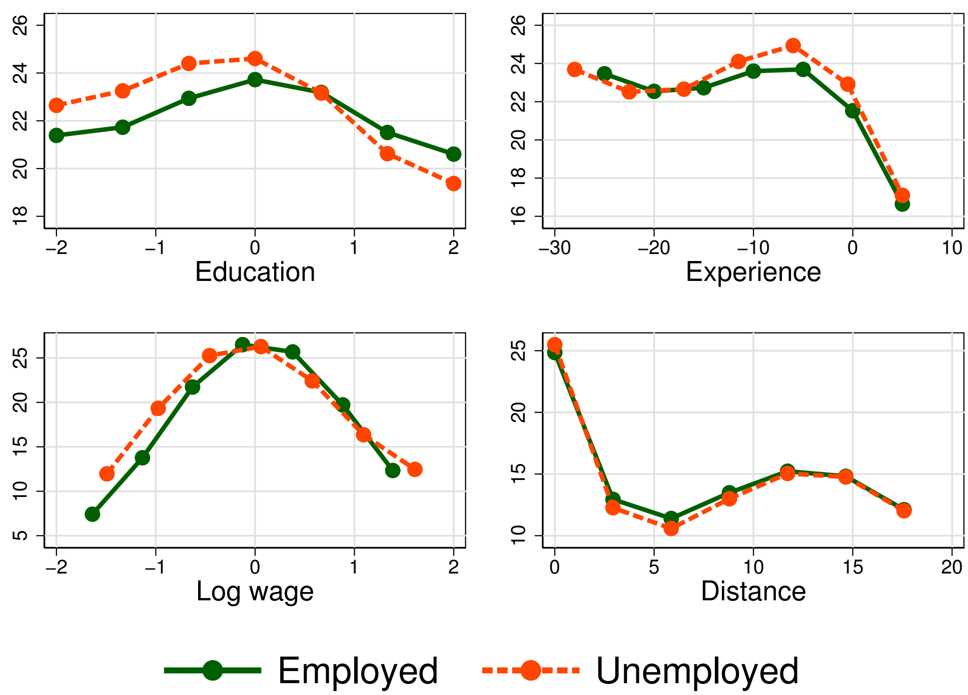

Figure 3: Predicted application probabilities, given results from eq. (3) and different levels of misalignment in the selected variable x (see main text for details). The rest of regressors are at their sample means.

In figure 3 we present graphically results of the effect of misalignment in years of education, years of experience, log wages, and regional distance on application decisions. The figure shows predicted application probabilities ($\hat{y}_{aw}$ from the estimates of equation 3), when a particular continuous dimension misalignment (z_k) varies, keeping all other observables at their sample mean, including the misalignment in other dimensions. Given that each misalignment dimension enters the equation as a fifth-order polynomial and that there are interactions between them, the computed effect is potentially highly non-linear and depends on which value the other control variables take. The considered range for z_k is bounded by its 5th and 95th percentiles.

As seen in figure 3, job seekers in both labor market states tend to align themselves with the advertised requirements of job postings. This is represented by an inverted U-shaped relationship between *misalignment* and application probability (all else constant) for education, experience, log wages, and by a mostly decreasing line in the case regional distance.

In the upper-left panel of figure 3, the application probability for both employed and unemployed peaks at zero, e.g., an exact alignment between required and realized years of education. Nevertheless, the shapes of employed and unemployed look asymmetrical. The application probability is larger for the unemployed when the applicant is overqualified, whereas the pattern reverses when the applicant is underqualified. Therefore, employed seekers seem less reluctant to apply to jobs for which they are underqualified in terms of education. In other words, employed workers are more ambitious or daring to take the next rung of the job ladder, assuming that jobs requiring more education are better.

For the unemployed, the experience dimension curves peak around -7 and show a steeper decline to the right. This means that job seekers tend to have more than seven years of experience than the minimum required and do not refrain from applying if they are even more overqualified in experience. The main reason for the average misalignment in this dimension is that most of our sample are individuals with a significant number of years of experience. The application probability curve for the employed peaks a bit to the left of the one for the unemployed, and the gap between the two groups becomes wider for experience gaps between -15 and 0, suggesting that the unemployed are slightly proner to apply to jobs for which they are overqualified in experience.

The plot in the lower-left panel reveals that differences in log wages greatly affect application probabilities: the application probabilities for the unemployed fluctuate between 12% and 25% for , while for employed seekers, the range is wider, from around 7% to 25%. Given that our estimates control for all other observables across job positions and job seekers and that the regression controls for interactions, we can interpret the misalignment in $\log$-wages as a gap in job and worker unobserved productivities. Controlling for all observables, higher-paying jobs and job seekers with higher earnings expectations must have higher skill levels on average, and vice versa. These interpretations align with our findings on high positive assortative matching at the application stage (Banfi, Choi, and Villena-Roldán, 2022). Overall, the unemployed curve surpasses to its employed counterpart for negative gaps. The unemployed are more likely to apply to jobs for which they are overqualified in terms of productivity. For jobs for which applicants are underqualified, to the right of the peaks, employed applicants apply more often. These facts portray on-the-job searchers as more daring than their unemployed counterparts, probably due to the better outside options of the former.

The lower-right panel depicts the predicted probability as a function of the distance between the regional capital of the applicant and the regional capital of the job in hundreds of kilometers. For ads located relatively close to the applicants, the likelihood of an application decreases quite quickly: from 25% at zero-distance to nearly 10% for a 600 kilometers distance. For higher distances there is some increase, which may be related to the geography of the country: because approximately 77% of the population lives less than 600 kilometers away from Santiago, which in itself represents 40%, individuals from north and south extremes of the country may have internalized moving to the central part of the country for better labor outcomes. The employed-unemployed gap is only noticeable in mid-distance applications. Marinescu and Rathelot (2018) and Manning and Petrongolo (2017) estimates imply a much larger drop in the likelihood of applying to jobs as distance increments, although our estimates are not directly comparable because we measure an intention rather than an effective reallocation and control for a substantially richer set of variables.

5 Conclusions

Using data from a Chilean job posting website, we uncover several facts regarding the nature of job search online. Given our unique setup, we can deconstruct behavior into two dimensions: *intensive* and *selective* margins. The first one refers to deciding the probability of application into a consideration set, i.e. an individually local labor market, while the second one concerns with the application decision responding to the potential match fit. Even though there are a number of papers trying to measure the intensity of job search in different ways, there are no other study documenting the kind of jobs workers apply to or forgo, to the best of our knowledge. To describe these decisions, we focus in the concept of misalignment , that is, the gap between a job ad requirement and the relevant characteristic of the worker. The richness of our database allows us to estimate the behavior of job seekers when facing misalignment in education, experience, log wages, geographical distance, and occupation. We also focus on differential behavior patterns by employment status.

While the estimation of the effects of different ad and worker characteristics and their interactions is relatively simple, determining the set of relevant ads for each applicant is not. Hence, we define a subset of ads that are similar to those workers actually apply. Instead of using a set of predetermined relevant dimensions to define the relevant submarket for the applicant, we use the bipartite network of applications and define a Jaccard metric between applicants which transpires into a proximity measure between applicants and ads. Besides this application, our network approach can be useful for industrial organization or marketing applications in which buyers face a large number of options but only a subset is relevant. For the intensive margin, we find that the unemployed apply a bit more than the employed. We also find that males apply more, especially if unemployed. Employed married apply less than their single counterparts. We also find that the elapsed job ad posting negatively affect the likelihood of application for the unemployed, but not for the employed. This partially supports stock-flow matching or phantom vacancy theories.

The intensive margin increases with age for the employed, perhaps due to an effort substitution effect given scarcer job opportunities for older workers or a self-selection into the job board among potential users. A mild decrease over the life cycle occurs to the unemployed. We also find decreasing effort as the job search duration increases, but a mild positive effect for tenure. The cyclicality of the intensive margin is nuanced: it is quite insensitive for unemployment rate below the median, but highly procyclical in regional/time markets of high unemployment. This may help explain disparate findings in the literature.

The selective margin matters for job search behavior. All workers negatively react to misalignment in terms of education, experience, log wages, distance, and occupation. Our analysis shows that employed job seekers are more *ambitious*: they tend to apply to jobs requiring more education than they have, and to job ads with wages above their own expectations. Log wage misalignment generates large changes in application probability.

Taking stock, we present a nuanced and general picture of job seeking behavior in dimensions rarely, if ever, studied before. We hope our findings could motivate further empirical research and help construct and discipline theoretical models.

References

ABALUCK, J., AND A. ADAMS-PRASSL (2021): "What do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses*," *The Quarterly Journal of Economics*, 136(3), 1611–1663.

AHN, H. J., AND L. SHAO (2017): "Precautionary on-the-job search over the business cycle," *Available at SSRN 2897533*.

ALBRECHT, J., B. DECREUSE, AND S. VROMAN (2023): "Directed search with phantom vacancies," *International Economic Review*, Forthcoming.

BAKER, S. R., AND A. FRADKIN (2017): "The impact of unemployment insurance on job search: Evidence from Google search data," *Review of Economics and Statistics*, 99(5), 756–768.

BANFI, S., S. CHOI, AND B. VILLENA-ROLDÁN (2022): "Sorting on-line and on-time," *European Economic Review*, 146, 104128.

BANFI, S., AND B. VILLENA-ROLDÁN (2019): "Do High-Wage Jobs Attract More Applicants? Directed Search Evidence from the Online Labor Market," *Journal of Labor Economics*, 37(3), 715–746.

BRANSCH, F. (2021): "Job search intensity of unemployed workers and the business cycle," *Economics Letters*, 205, 109927.

CARRILLO-TUDELA, C., AND L. VISSCHERS (2023): "Unemployment and Endogenous Reallocation over the Business Cycle," *Econometrica*, Forthcoming.

CHÉRON, A., AND B. DECREUSE (2016): "Matching with Phantoms," *Review of Economic Studies*, (0), 1–30.

CHOI, S., A. JANIAK, AND B. VILLENA-ROLDÁN (2015): "Unemployment, Participation and Worker Flows Over the Life-Cycle," *The Economic Journal*, 125(589), 1705–1733.

COLES, M., AND B. PETRONGOLO (2008): "A Test between Stock-Flow Matching and the Random Matching Function Approach," *International Economic Review*, 49(4), 1113–1141.

COLES, M. G., AND A. MUTHOO (1998): "Strategic Bargaining and Competitive Bidding in a Dynamic Market Equilibrium," *Review of Economic Studies*, 65(2), 235–260.

COLES, M. G., AND E. SMITH (1998): "Marketplaces and Matching," *International Economic Review*, 39(1), 239–254.

DAVIS, S. J., AND B. SAMANIEGO DE LA PARRA (2017): "Application Flows," Unpublished manuscript.

DELLAVIGNA, S., J. HEINING, J. F. SCHMIEDER, AND S. TRENKLE (2021): "Evidence on Job Search Models from a Survey of Unemployed Workers in Germany," *The Quarterly Journal of Economics*, 137(2), 1181–1232.

DELOACH, S. B., AND M. KURT (2013): "Discouraging workers: estimating the impacts of macroeconomic shocks on the search intensity of the unemployed," *Journal of Labor research*, 34, 433–454.

DINARDO, J., N. M. FORTIN, AND T. LEMIEUX (1996): "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica*, 64(5), 1001–1044.

EBRAHIMY, E., AND R. SHIMER (2010): "Stock-flow matching," *Journal of Economic Theory*, 145(4), 1325–1353.

FABERMAN, R. J., AND M. KUDLYAK (2019): "The intensity of job search and search duration," *American Economic Journal: Macroeconomics*, 11(3), 327–57.

FLUCHTMANN, J., A. GLENNY, N. HARMON, AND J. MAIBOM (2023): "Unemployed Job Search across People and over Time: Evidence from Applied-for Jobs," *Journal of Labor Economics*, forthcoming.

FOOTE, A., M. J. KUTZBACH, AND L. VILHUBER (2021): "Recalculating...: How Uncertainty in Local Labour Market Definitions Affects Empirical Findings," *Applied Economics*, 53(14), 1598–1612.

FUJITA, S. (2012): "An empirical analysis of on-the-job search and job-to-job transitions," Working paper, Federal Reserve Bank of Philadelphia.

GOMME, P., AND D. LKHAGVASUREN (2015): "Worker search effort as an amplification mechanism," *Journal of Monetary Economics*, 75, 106–122.

GREGG, P., AND B. PETRONGOLO (2005): "Stock-flow matching and the performance of the labor market," *European Economic Review*, 49(8), 1987–2011.

HAZELL, J., AND B. TASKA (2020): "Downward rigidity in the wage for new hires," Discussion paper, SSRN 3728939.

HERZ, B., AND T. VAN RENS (2020): "Accounting for mismatch unemployment," *Journal of the European Economic Association*, 18(4), 1619–1654.

HONKA, E., A. HORTAÇSU, AND M. WILDENBEEST (2019): "Empirical search and consideration sets," vol. 1 of *Handbook of the Economics of Marketing*, pp. 193–257. North-Holland.

HORNSTEIN, A., P. KRUSELL, AND G. L. VIOLANTE (2011): "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," *The American Economic Review*, 101(7).

JACCARD, P. (1901): "Étude de la distribution florale dans une portion des Alpes et du Jura," *Bulletin de la Societe Vaudoise des Sciences Naturelles*, 37, 547–579.

JOVANOVIC, B. (1979): "Job matching and the theory of turnover," *Journal of political economy*, 87(5, Part 1), 972–990.

KUDLYAK, M., D. LKHAGVASUREN, AND R. SYSUYEV (2013): "Systematic Job Search: New Evidence from Individual Job Application Data," mimeo, Federal Reserve Bank of Richmond.

LEYVA, G. (2018): "Against All Odds: Job Search During the Great Recession," Unpublished Manuscript. Banco de Mexico.

LI, F., AND X. WENG (2017): "Efficient Learning and Job Turnover in the Labor Market," *International Economic Review*, 58(3), 727–750.

MANNING, A., AND B. PETRONGOLO (2017): "How local are labor markets? Evidence from a spatial job search model," *American Economic Review*, 107(10), 2877–2907.

MARINESCU, I., AND R. RATHELOT (2018): "Mismatch unemployment and the geography of job search," *American Economic Journal: Macroeconomics*, 10(3), 42–70.

MARINESCU, I., AND R. WOLTHOFF (2020): "Opening the black box of the matching function: The power of words," *Journal of Labor Economics*, 38(2), 535–568.

MENZIO, G., I. A. TELYUKOVA, AND L. VISSCHERS (2016): "Directed search over the life cycle," *Review of Economic Dynamics*, 19, 38 – 62, Special Issue in Honor of Dale Mortensen.

MICHELACCI, C., AND J. SUAREZ (2006): "Incomplete Wage Posting," *Journal of Political Economy*, 114(6), 1098–1123.

MORETTI, E. (2011): "Local Labor Markets," vol. 4 of *Handbook of Labor Economics*, pp. 1237–1313. Elsevier.

MUKOYAMA, T., C. PATTERSON, AND A. ŞAHIN (2018): "Job Search Behavior over the Business Cycle," *American Economic Journal: Macroeconomics*, 10(1), 190–215.

NAUDON, A., AND A. PÉREZ (2018): "Unemployment dynamics in Chile: 1960-2015," *Journal Economía Chilena (The Chilean Economy)*, 21(1), 4–33.

PISSARIDES, C. A., AND J. WADSWORTH (1994): "On-the-job search: Some empirical evidence from Britain," *European Economic Review*, 38(2), 385–401.

ŞAHIN, A., J. SONG, G. TOPA, AND G. L. VIOLANTE (2014): "Mismatch unemployment," *American Economic Review*, 104(11), 3529–64.

TAYLOR, C. R. (1995): "The Long Side of the Market and the Short End of the Stick: Bargaining Power and Price Formation in Buyers', Sellers', and Balanced Markets," *The Quarterly Journal of Economics*, 110(3), 837–855.

VAN NIEROP, E., B. BRONNENBERG, R. PAAP, M. WEDEL, AND P. H. FRANSES (2010): "Retrieving unobserved consideration sets from household panel data," *Journal of Marketing Research*, 47(1), 63–74.

VAN OMMEREN, J., AND G. RUSSO (2013): "Firm Recruitment Behaviour: Sequential or Non-sequential Search?," *Oxford Bulletin of Economics and Statistics*, pp. 1–24.

VAN OURS, J., AND G. RIDDER (1992): "Vacancies and the Recruitment of New Employees," *Journal of Labor Economics*, 10(2), 138–155.

Deconstructing Job Search Behavior

Online Appendix

not intended for publication

Table A1: Intensive margin coefficients by labor status, no compositional adjustment

VARIABLES	(1) Employed weights	(2) Employed no weights	(3) Unemployed weights	(4) Unemployed no weights
Married	-2.249***	-0.469	0.513	-0.030
	(0.628)	(0.322)	(0.363)	(0.179)
Male	1.306***	0.203***	2.122***	0.314***
	(0.055)	(0.028)	(0.044)	(0.022)
Explicit wage (w)	0.097**	-0.121***	0.604***	-0.060***
	(0.049)	(0.025)	(0.041)	(0.021)
Explicit wage (a)	-2.734***	-0.682***	-0.236***	0.030
	(0.078)	(0.039)	(0.056)	(0.029)
No. of Vacancies (a)	-0.031***	-0.004**	0.023***	0.011***
	(0.004)	(0.002)	(0.002)	(0.001)
Ad duration (weeks)	-0.007***	0.002*	-0.139***	-0.024***
	(0.002)	(0.001)	(0.002)	(0.001)
Observations	2,896,402	2,896,402	4,189,865	4,189,865
R-squared	0.129	0.038	0.119	0.037
Mean app prob	25.43	4.679	26.07	4.830

Notes: Regression coefficients from a linear regression on application decisions. Dependent variable is y_{aw}, a dummy for the existence of a job application. Each regression controls also for polynomials and interactions in *misalignment* as well as age of the worker, firm size, contract type, dummies for different types of requirements of the job and characteristics of the firm (see details in the main text). Standard errors in parentheses. One, two, and three asterisks indicate significance at 10%,5%, and 1%, respectively.

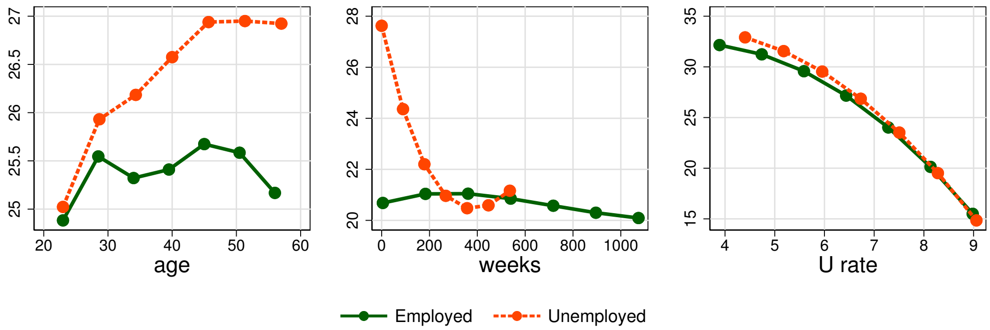

Figure A1: Predicted application probabilities for different ages, number of weeks in the current labor force status, and national unemployment rate at the time of the application decision, given results from equation (3), no compositional adjustment. The figure is computed using the coefficients associated to a polynomial of order 5 on each variable and leaving the rest of regressors at their sample mean.

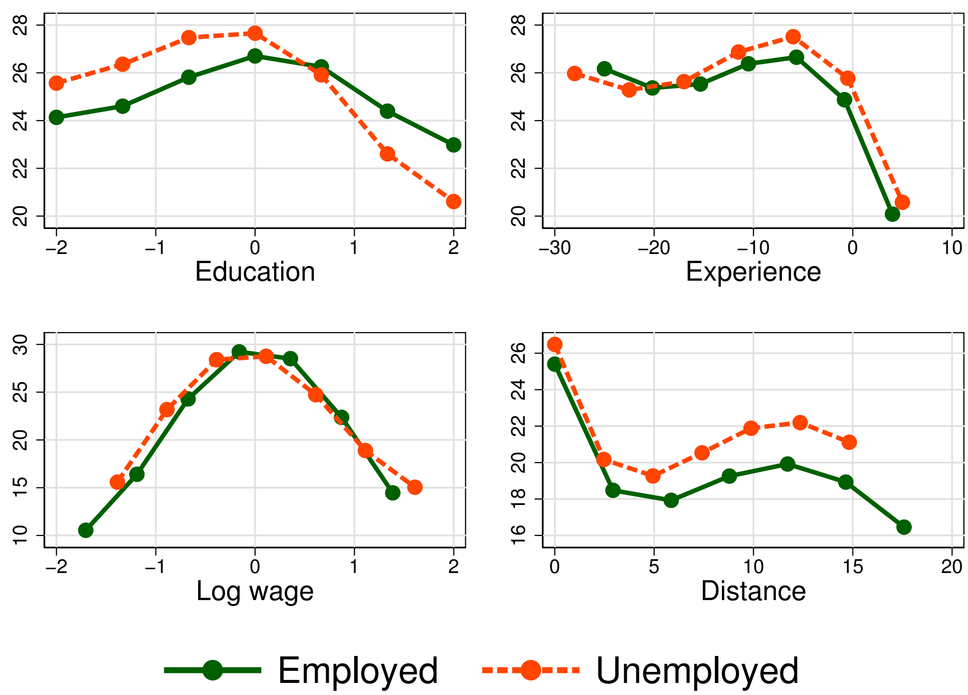

Figure A2: Predicted application probabilities, given results from eq. (3) and different levels of misalignment in the selected variable x (see main text for details). The rest of regressors are at their sample means. no compositional adjustments