Deconstructing Job Search Behavior*

Stefano Banfi
Ministry of Energy, Chile

Sekyu Choi
University of Bristol

Benjamín Villena-Roldán
CEA, DII, University of Chile,
SMAUG, MIPP

June 16, 2018

Abstract

We use an unusually rich data from a Chilean job board to document various theoretically relevant facts regarding job search. We show how application behavior is influenced by (1) demographics such as gender, age, and marital status, (2) alignment between applicant wage expectations and wage offers, (3) applicant fit into ad requirements such as education, experience, job location and occupation (4) timing variables, including unemployment duration, job tenure (for on-the-job searchers) and business cycle conditions. Our paper provides novel evidence that can discipline current and future search-theoretical frameworks.

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

JEL Codes: E24, J40, J64

*Email: stefano.banfi@gmail.com, sekyu.choi@bristol.ac.uk and bvillena@dii.uchile.cl. 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, and Universidad de Chile for insightful discussion and comments. All errors are ours.

Introduction

What kind of jobs workers look for and how much effort they exert are critical pieces of information for labor market outcomes. These search activities determine wages and job allocations in the economy to a large extent. Key magnitudes for contemporaneous policy debates such as the unemployment rate, mobility, and income inequality are affected by jobseekers' decisions. We economists have some evidence about *extensive* and *intensive margins* of search effort, but we know very little about the kind of jobs workers search for, what we call a *selective margin*. As applications are tentative allocations, they are of prime importance to understand ex post outcomes. Hence, while we keep an eye on the search effort, our main contribution is to carefully document nuanced application decisions workers make according to their own characteristics, job requirements, and context, and discuss their implications.

Data from online job boards offer us the opportunity to learn about job search more deeply than ever before. While online job search is different from traditional methods, this modality deserves increasing attention as its importance and efficiency has increased over time (Kuhn and Mansour, 2014). In most theoretical models, job seeking behavior is highly stylized: workers accept or reject offers using simple optimal rules. However, jobs are complex objects with many interlinked aspects which matter for job seekers: wages, fit, location, timing, etc. On the other hand, job seekers' characteristics may also affect optimal search strategies.

We use information from `www.trabajando.com`, a job posting website with presence in most of Latin America, in addition to Spain and Portugal. We exploit a comprehensive dataset on daily applications of job seekers to job postings in the Chilean labor market during the period 2008 to 2016. Our dataset contains some novel features. We observe detailed information on both sides of the market: we observe education, occupations,[1] and experience for individuals and for job postings (as requirements stipulated by firms). Moreover, we observe detailed information of seeker and job characteristics, as well as both desired and current wages for individuals (wages of last full time jobs if unemployed) and the wages firms expect to pay at jobs they are posting.

The main research question in this paper is how individuals, facing a set of online job ads, choose to apply to some jobs and forgo others. Given our unique setup, we can deconstruct behavior into two dimensions: an intensive dimension, the number of applications sent, which relates to standard notions in the literature,[2] and a selective one, where we analyze what affects the decision of sending an application at the margin. For the latter, we estimate application decision equations, in which we disentangle the contribution of a large array of factors influencing application choices.

In this paper we also provide a methodological contribution. To overcome the fact that we only observe effective applications and not the entire set of relevant job positions for each candidate,[3] we

[1] We observe a one-digit classification, created by the website administrators.

[2] See Mukoyama, Patterson, and Şahin (2018) for example.

[3] Not observing page views in the website is a shortcoming in the literature using online job board data. Unfor-

use the bipartite *network* formed by job seekers linked through common job applications. Using this network, we construct the choice set of an applicant as the list of all ads applied by seekers linked with her, and create weights given a similarity metric between two job seekers. Our methodology overcomes the problem of only observing effective applications in the original data.

Instead of this network approach, 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 (2015), or by clustering algorithms as Banfi and Villena-Roldan (2016). Our procedure has two important advantages. 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 often few arbitrary dimensions such as locations, occupations, 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 (2014). Second, the network approach allows us to define individual choice sets, generating key variability for parameter identification.

Our empirical exercise reveals interesting patterns with respect to job seekers' application decision making and search effort, measured by the probability of applying to a position in the relevant set. We find that some demographic characteristics are quite relevant at the margin: men apply to more job ads than women, while married individuals apply more than their single counterparts if they are unemployed. We also document that search effort decreases with age, which is consistent with evidence in Choi, Janiak, and Villena-Roldán (2015) and Menzio, Telyukova, and Visschers (2016), among others. This decrease in effort is of crucial importance for unemployment insurance design, as studied in Michelacci and Ruffo (2015). In terms of aggregates, we find that the number of submitted applications is counter cyclical, but the marginal application decision is affected in a non-linear way by aggregate conditions.

One of our main findings relates to how job seekers align themselves with heterogeneous types of jobs. We find that search behavior is highly sensitive to the requirements of educational level, experience, location and occupation requirements and that job seekers target an optimal or most preferred type of job, which is not necessarily the one that matches perfectly their current characteristics: the probability of an application peaks when the applicant is slightly underqualified in terms of education but the pattern is reversed in the case of experience requirements. An implication of this result, is that job seekers' incentives to apply to jobs that are further away from the most preferred job (in terms of these characteristics) decrease, although this decrease is not symmetrical. We also study how this fit evolves over the life-cycle, the duration in the current labor force status and across different levels of aggregate unemployment rates (business cycle conditions).

tunately, `www.trabajando.com` nor other job search boards keep records of page views by applicant for two reasons: (i) it is very expensive to keep these records while the information is of little use (for the job board operators), and (ii) applicants need to be logged in when viewing job ads, a requirement that would reduce the likelihood of getting applicants into the board. See the references below.

In terms of wages, we find that individuals are more likely to apply for a job offering a wage close to their expectations. Going one step further, and given that in our empirical approach we control for an exhaustive array of observable characteristics, one could argue that wages in our setup are proxies for inherent *types* of both workers and firms. Thus, our results are suggestive of positive assortative matching patterns at the application stage, but more importantly, summarizes well the main distinction between search strategies of the unemployed versus the employed: while the unemployed target wage offers (types) that are very close to their own stated salary expectations, employed seekers target wages (types) which are on average above their expectations. Our reading from this is that the unemployed are trying to maximize the chances of obtaining a job offer, while the employed (performing on-the-job search) are more likely trying to climb the job ladder.

Our results also show that application decisions decline with either unemployment duration (for unemployed seekers) and job tenure (for those performing on-the-job search). This evidence is particularly useful to understand the dynamic evolution of unemployed workers over an unemployment spell, an important input for the design of unemployment insurance policies, an aspect also considered by Faberman and Kudlyak (2013); the effect of tenure on job search is also relevant to understand factors behind job-to-job transitions, arguably an important mechanism to explain wage dispersion, as stated in Hornstein, Krusell, and Violante (2011)).

Our paper is related to a growing literature which use data from online job-posting/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 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). Marinescu and Rathelot (2015) 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 (2015) 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 of explicit wages. Gee (2015), using a large field experiment on the job posting website `www.linkedin.com`, shows that being made aware of the number of applicants for a job, increases ones own likelihood of making application.

On product markets, Lewis (2011) shows that internet seekers for used cars significantly react to posted information regarding automobiles' quality. Jolivet and Turon (2014) and Jolivet, Jullien, and Postel-Vinay (2016) use information from a major French e-commerce platform, `www.PriceMinister.com`, to study the effects of search costs and reputational issues (respectively) in product markets.

The data

We use data from `www.trabajando.com` (henceforth the website) a job search engine with presence in mostly Spanish speaking countries: as of September of 2017, the list comprises Argentina, Brazil, Colombia, Chile, Mexico, Peru, Portugal, Puerto Rico, Spain, Uruguay and Venezuela. Our data covers a sample of job postings and job seekers in the Chilean labor market, 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.[4] We have codes for occupational area of the current/last job of individuals, 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.[5]

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, an occupational code, geographic information (”región” only) and some limited information on the firm offering the job: its size (number of employees in brackets) and industry. Educational categories are *primary* (one to eight years of schooling), *high school* (completed high school diploma), *technical tertiary education* (professional training after high school), *college* (completed university degree) and *post-graduate* (any schooling higher than 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 given the same choice whether to make this information visible or not to the applicants. Naturally, one could question the reliability of wage information which will be ultimately hidden from the other side of the market. Banfi and Villena-Roldán (2016) 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

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

[5]As of January 3rd, 2018, the 60-day fee is CLP 69,900 + 19% VAT as posted in `http://www1.trabajando.cl/empresas/noticia.cfm?noticiaid=3877`, which is equivalent to USD 136 or EUR 113.

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: individuals seeking for jobs through other means, and more importantly, outcomes of job applications.

For the remainder of the paper, we restrict our sample to consider only individuals working under full-time contracts and those 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 represents much more than the 99th percentile of the 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 unit 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) in the website and the time of their last submitted application 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 further drop individuals who apply to more than the 99-th percentile of job applicants in terms of number of submitted applications.

Table 1 shows descriptive statistics for the job searchers 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 ($86,687$ 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, according to the 2013 CASEN survey. There is also a big discrepancy by labor force status: unemployed seekers are significantly less educated

[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%), to disability insurance (1.2%), and mandatory contribution to unemployment accounts(0.6%)

[7] Using the average nominal exchanges rate between 2013-16, https://si3.bcentral.cl/Siete/secure/cuadros/home.aspx.

[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 http://observatorio.ministeriodesarrollosocial.gob.cl/casen-multidimensional/casen/casen_2015.php

Table 1: Characteristics of Job Seekers

	Employed	Unemployed	Total
Demographics			
Age	33.77	33.25	33.46
Males	0.62	0.54	0.57
Married	0.34	0.28	0.30
Experience (years)	8.28	7.64	7.90
Wages (thousand CLP)	1,087	592	792
Tenure (weeks)	177.96	–	177.96
Unemployment duration (weeks)	–	60.20	60.20
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	4.83	4.99
Number of applications	1.49	1.53	1.52
Observations	86,687	128,482	215,169

in 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 searching for a job is higher for those employed than for the unemployed: 5.24 versus 4.83 weeks respectively. In terms of applications, both groups show very similar choices, with around 1.52 submitted applications.

Number of applications

In this section, we compute a regression relating the number of submitted applications by individuals, controlling for some individual characteristics and aggregate conditions. In table 2, we show results of a negative binomial regression, where we split the sample by employment status of the job seeker at the start of their search window.

[9]The minimum wage has increased substantially in recent years. For information about the trajectory of the legal minimum wage in Chile, please see https://www.leychile.cl/Consulta/listado_n_sel?_grupo_aporte&sub=807&agr=2.

Table 2: NB regression on number of submitted applications

	Unemployed	Employed
married	0.0700	0.0200
	(0.443)	(0.919)
male	0.0488***	0.0150
	(0.000)	(0.500)
married × male	-0.0229	0.0264
	(0.342)	(0.497)
explicit wage worker	0.0164	0.0026
	(0.126)	(0.884)
Weeks searching	0.0084***	0.0494***
	(0.000)	(0.000)
pseudo R^2	0.014	0.010
N	10,097	25,881

Notes: Regression coefficients from a Negative Binomial regression on the number of applications submitted. P-values in parenthesis. Each regression includes a fifth order polynomials on the age of the worker, tenure/unemployment duration (in weeks) and the aggregate unemployment rate in the economy.

In the regression we also control (not shown in the table) for categorical variables related to the region of residence of the seeker, a one-digit occupational code, education attainment, years of experience and fifth order polynomials on age, number of weeks in the current labor market state (employment or unemployment) and the aggregate unemployment rate in the Chilean economy.

As seen from the table, the male dummy and the number of weeks searching on the website have significant effects on the number of submitted applications. The latter effect is probably the mechanical effect of more submitted applications being correlated with longer search windows by individual.

On the other hand, in figure 1 we plot the predictions of the regression model when we vary age (left panel), the number of weeks in the current labor force status (center panel) or the aggregate unemployment rate (right panel) in a range comprised the sample mean of the variables, plus and minus 1.5 times their respective standard deviation. The predictions are computed leaving the rest of explanatory variables in the regression of table 2 at their sample mean.

There are several interesting takeaways from this initial analysis. Although only significant for the Unemployed sample, male seekers seem to put more effort in terms of number of applications sent. This is not however, correlated with marital status. In table 2, we also see that the coefficient related to the total number of weeks searching is positive and significant for both groups. This is primarily a mechanical result: weeks of searching is determined by the total number of weeks

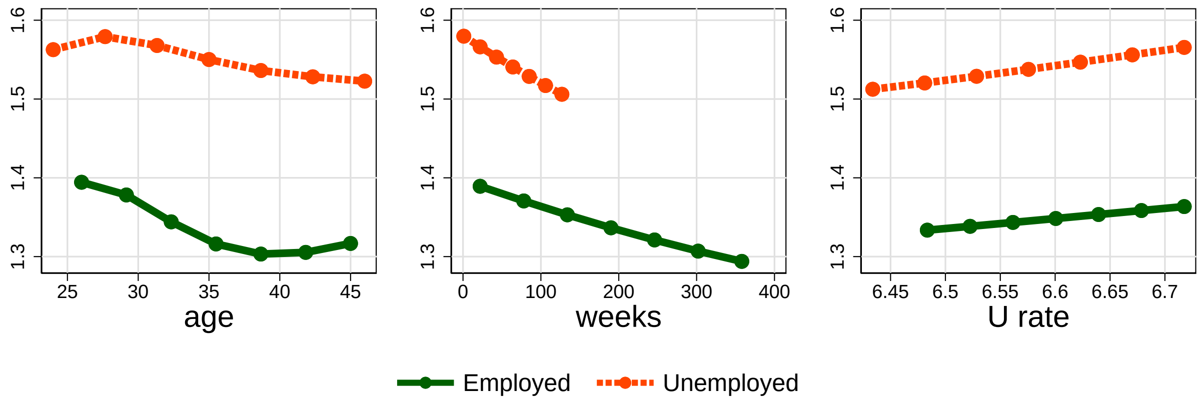

Figure 1: Predicted number of submitted applications for different ages, given results from table 2. The figure is computed using the coefficients associated to a polynomial of order 5 on age of the applicant.

between the first and last applications: thus, there is an obvious correlation between longer time search windows and number of applications sent. This is specially true since a large fractions of individuals send only one application.

There are three main results in figure 1: first, there are life-cycle effects in search effort, which is reflected in decreasing number of predicted applications sent for both employed and unemployed in the left panel, with the unemployed exhibiting a small hump shaped curve peaking at age 27 approximately. Second, we find a time effect in the center panel, in that the number of applications declines with the number of weeks job seekers have been in either unemployment or in their current job. Third and final, the right panel shows that search efforts (in terms of total number of applications sent) is counter cyclical, since predicted numbers increase with the level of total unemployment rate in the economy.

Application probabilities and job seeker preferences

In this section we analyze empirically which attributes of heterogeneous jobs are more valued by heterogeneous job seekers. To do this, we first need to determine which is the relevant market segment for each individual in our sample. However, our dataset only contains information on actual applications and no information is collected by the website on total number of searches nor *clicks* on job postings by individuals. Thus, we do not have sample variation in terms of job ads: we only observe those that individuals chose to apply to, but not those which are observed but then discarded by seekers.

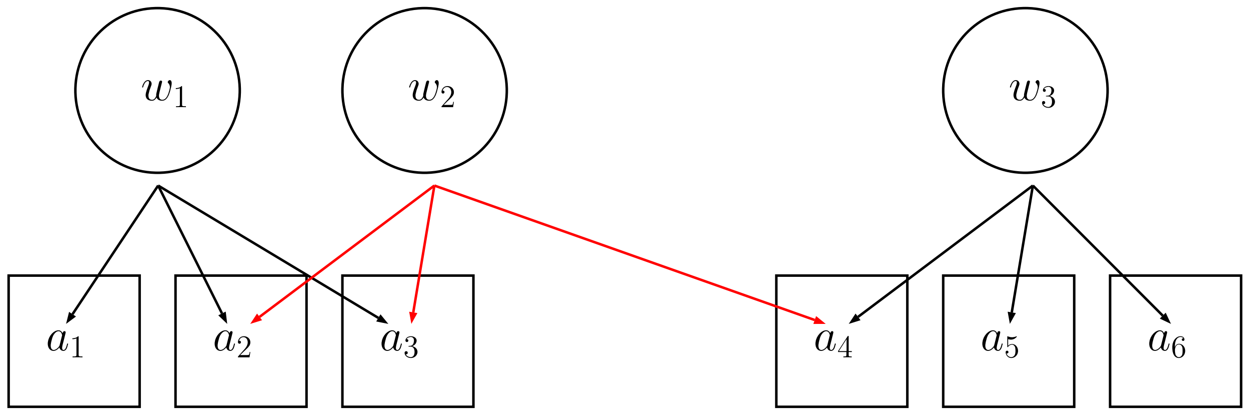

Figure 2: Example of a network formed by workers $\{w_1, w_2, w_3\}$. Worker w_1 is linked to worker w_2 by common applications to ads a_2 and a_3 but is not linked with w_3 in the network of degree 1. All workers are linked in the network of degree 2.

Market segmentation through network analysis.

One way to obtain the variation described above would be to consider the cross between all job seekers and all job ads in our sample, what is usually referred to as the *exploded* dataset. However, there are some drawbacks from this approach: first, the *exploded* dataset makes comparisons between job seekers and job positions which may be objectively too different to consider; second, the size of the estimating sample increases exponentially with the initial number of job seekers and job ads, making the task of even simple calculations infeasible in most cases.

In what follows, 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 can 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 of 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 which 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 dataset $\mathcal{A}_w^2$. We can continue with this logic iteratively, until forming the set $\mathcal{A}_w^\infty$, which is the cross between each job seeker w and all job postings a (or the *exploded* dataset).

Figure 2 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 characteristic of a_4 is 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 $\mathcal{A}^0_{w_1} = \{w_1, w_2, w_3\}$, and as discussed above, $\mathcal{A}^1_{w_1} = \{a_1, a_2, a_3, a_4\}$. 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

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

which can be generalized for other degrees of separation.[10] In what follows, we will concentrate on networks of degree one only.

Table 3: Number of relevant ads (a) per worker (w)

	Potential ads for a worker			Potential workers for an ad		
	All	U	E	All	U	E
percentile 10	2	2	2	1	2	3
percentile 50	16	16	19	21	25	33
percentile 90	96	104	87	233	325	385
mean	38.5	40.7	36.8	96.5	123.8	147.8
standard deviation	68.1	73.8	57.1	278.7	342.0	410.3

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).

In table 3, 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 postings per job seeker, with employed seekers being related to more posts (19) than those unemployed (16). The number of potential ads exhibits quite the amount of variation, going from 2 (tenth percentile of distribution) to 104 and 89 for unemployed and employed respectively (ninetieth percentile). On the other hand, the median number of workers related to each job advert is 233, with employed individuals being attached to more job ads than those

[10] The generalization follows a recursive definition

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

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

unemployed. Given the sets of related job ads, mean application rates,[11] 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.

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 2. Since w_2 and w_1 submit common applications to several common positions, they must have similar preferences and qualifications. Then, the likelihood that w_2 truly considers applying to job ads to which w_1 applied to must be high. In contrast, w_2 and worker w_3 share less 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 $q(w, a)$ for each worker w and job position a. We start with 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 impose function $b(w, v)$ to be monotonic in set similarity.[12] A particular functional form that satisfies these conditions is

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

where $|S|$ is the cardinality (number of elements) of set S. Equation (1) is also known as the *Jaccard Similarity Index* between two groups.[13]

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

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

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,

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

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

[12] Note that these conditions are similar to conditions used to define distance in metric spaces.

[13] See Jaccard (1901).

linked to w. It is easily verified that given this proposed weighting function:

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

Preferences over heterogeneous characteristics.

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

$$y_{aw} = X_{aw}\beta_{aw} + \sum_{\{k_c,k_d\}} \sum_{p=1}^{\bar{P}} \{\beta_{k_o p}(z_{k_o})^p\} + +\beta_{k_o} z_{k_o} + \sum_k \sum_\ell \mathbf{1}_{\{k \neq \ell\}} \beta_{k\ell} z_k z_\ell + \epsilon_{aw} \tag{3}$$

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 control for observed job and worker characteristics, which do not overlap. The list of variables for the job includes firm's size, dummies for firm's industry (1 digit) and specific job requirements (computer knowledge, or some other form of specific knowledge) and controls for specific job characteristics: type of contract (full/part time), number of vacancies needed to be filled and controls for job titles.[14] 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 for the amount of time (measured in weeks) in either the current job (for those employed) or in unemployment (for unemployed seekers). For both seekers and ads, we include a variable of whether the wage expectation (for seekers) or the wage expected to be paid (for jobs) is made explicit or not. As control for business cycle conditions, we include the national unemployment rate of the Chilean economy during the date (quarter) in which the application took place.[15]

On the other hand, we include a set of controls for the *misalignment* (which we denote by z) between characteristics required by firms vs. the characteristics of the job seeker. For continuous variables, which we denote by k_c, and encompass the level of education, years of experience and log wages, we define z_{k_c} as the simple difference between the value of the characteristic required by the position and value of the characteristic possessed by the job seeker. For location variables, we construct a coarse metric z_{k_d}: if the job position and the individual are located in the same

[14] We follow work by Marinescu and Wolthoff (2015) and Banfi and Villena-Roldán (2016).

[15] 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 tha particular job ad.

region, the distance is zero. Otherwise, z_{k_d} takes the absolute difference between regions. Given the special geography of Chile, one can order regions from north to south. Thus, although coarse, our measure z_{k_d} is informative of distances workers are willing to consider for jobs.

For occupations, the variable z_{k_o} is defined as a dummy that takes the value of one when the category in the job posting is different from the characteristic of the worker and zero when they are the same.

In equation (3), for each of the continuous dimensions k_c and for the geographic dimension k_d, 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. The basic idea is to try to understand if agents apply differently if they are *over*-qualified ($z_{k_c} < 0$) compared to when they are *under*-qualified ($z_{k_c} > 0$) or if there are non-linear effects of distance on considered application decisions (variable k_d. We estimate the above equation separating our sample between the employed and unemployed, in order to assess whether on-the-job search differs from unemployed search behaviour. Finally, we also consider interaction effects between different *misalignment* levels and weight worker-ad observations by function q described above.

Results on individual application decisions

Table 4 shows results from estimating equation (3) using ordinary least squares. The table shows coefficients multiplied by 100, separates by employment status and whether we perform weighting or not of our estimates (according to the discussion above). Results related to polynomials on continuous *misalignment* variables are presented later. Notice that, as opposed to table 2 where only independent variables related to workers were included, now we show how job ads (a) characteristics and *misalignment* measures between job workers and ads affect individual application decisions.

In the table we see a number of interesting results. In terms of demographics and family composition, we find that married individuals tend to apply marginally to fewer positions, while males are more likely to do so. The effect of the interaction term of being married and a male job seeker is not clear: the coefficient is positive and significant for the unemployed, but negative and insignificant for the employed.

Our results show that individuals who choose to be explicit about their wage expectations at the time of an application are more likely to apply to a position, which is true for both unemployed and employed seekers, although the effect is stronger for the employed sample. As with the results in table 2, where we show regression coefficients for the number of total applications submitted, the time (in weeks) that job seekers have searched in the website is positively related to submitting additional applications.

For postings, the table shows that an explicit wage in the job ad affects negatively the decision to apply for for all type of job seekers, but the effect is less important for the employed. On the other

Table 4: Regression Results

	Unemployed		Employed	
	weighted	non weighted	weighted	non weighted
Married (w)	-0.321	-0.007	-4.913^{***}	-1.161^{*}
	(0.467)	(0.972)	(0.000)	(0.012)
Male (w)	0.731^{***}	0.007	0.412^{***}	-0.007
	(0.000)	(0.815)	(0.000)	(0.874)
Male x Married (w)	0.589^{***}	0.084	-0.0759	-0.001
	(0.000)	(0.150)	(0.642)	(0.989)
Explicit wage (w)	1.353^{***}	0.128^{***}	1.546^{***}	0.213^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
Worker search duration (weeks)	0.0838^{***}	-0.0104^{***}	0.0914^{***}	0.0271^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
Explicit wage (a)	-1.227^{***}	-0.230^{***}	-3.051^{***}	-0.712^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
No. of vacancies (a)	0.0469^{***}	0.011^{***}	-0.005	-0.000
	(0.000)	(0.000)	(0.358)	(0.952)
Time ad (a) available (weeks)	-0.196^{***}	-0.031^{***}	0.0253^{***}	0.009^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
Different occupation (1 digit)	-12.140^{***}	-3.755^{***}	-8.809^{***}	-2.751^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
Adjusted R^2	0.076	0.027	0.064	0.022
Sample size	$2,998,831$	$2,998,831$	$1,265,887$	$1,265,887$

Notes: Regression coefficients from a linear regression on application decisions. Dependent variable is y_{aw}, a dummy for the existence of a job application. P-values in parenthesis. Each regression controls also for polynomials and interactions in *misalignment* as well as age of the worker, firm size, contract type, dummmies for different types of requirements of the job and characteristics of the firm (see details in the main text).

hand, unemployed individuals react positively to job ads that advertise higher number of vacancies to be posted, while those employed do not care about that information, given the insignificance of the estimated coefficient in that group. Finally, the effect of the perceived "age" of the job ad has also different effects depending on the labor force status of the individual: unemployed seekers seem to have a distaste for job ads that are older (in weeks), while those employed prefer them at the margin.

Results when we do not weight application by function $q(w, a)$ are shown in columns three (for unemployed) and five (for employed) of table 4. With the sole exception of the coefficient associated to the time the job ad is available online having a different sign for the unemployed pool, all other estimates have the same sign as in the weighted regressions. The value of all estimated coefficients increases as well as the adjusted R-squared, whenever we consider weights. This shows that the relative similarity between workers when assigning job ads, has big implications for the estimation results.

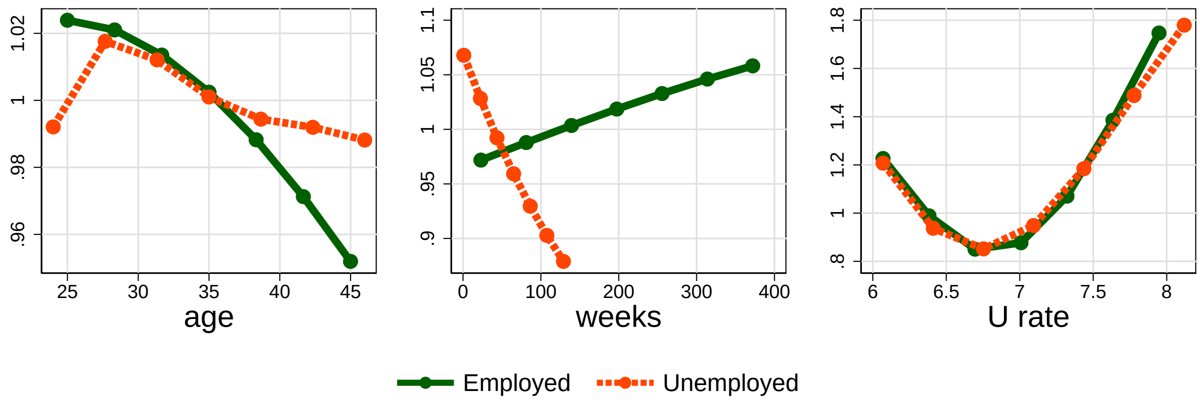

Figure 3: Predicted application probabilities (relative to sample averages) for different ages, number of weeks in the current labor force status and national unemployment rate at the time of the application decision, given results from eq. (3). The figure is computed using the coefficients associated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

Life-cycle, duration and business cycle effects.

Given the results from the regression, here we show how age (life-cycle effects), number of weeks in the current labor market status (duration dependence effects) and the aggregate unemployment rate at the time of application decisions (business cycle effects) affect application decisions. Given that these variables enter as polynomials of degree 5 in equation (3), in what follows we present results as predicted application probabilities relative to mean application rates in each group. While the methodology to obtain these figures is identical to the one used to produce figure 1, the difference is in the interpretation (total number of submitted applications versus individual worker-ad decisions) and how we control for both worker and ad characteristics in this section's results.

Figure 3 presents the results of the exercise. In the left panel, we observe that there are life-cycle effects in terms of job application decisions: for the employed, applications decrease with age, while for the unemployed, there is a hump-shaped profile. For both types of seekers, job applications decrease from early on (around age 27). This evidence is consistent with findings in Choi, Janiak, and Villena-Roldán (2015) and Menzio, Telyukova, and Visschers (2016), among others, with respect to job finding rates and employment to employment transitions over the life-cycle.

The middle panel in the figure shows that while unemployed have decreasing rates of applications as time passes (similar to what is found in figure 1), decreasing around 20 percentage points (from 1.05 to 0.85), the profile for those performing on-the-job search is mildly increasing, going from around 0.97 to around 1.05.

In terms of business cycle conditions, the right panel of figure 3 shows how application decisions react to the unemployment rate in the Chilean economy. The figure shows a nonlinear relationship between aggregate activity and application decisions, especially for low levels of unemployment. However, for the most part there seems to be a positive relationship between unemployment rates and applications, i.e., that search efforts are counter cyclical.

Misalignment and applications.

In table 4, the effect of *misalignment* in the occupation variable (z_{k_o}) can be read directly. The negative value associated to differences in occupational group between individual and job ad indicates that job seekers align themselves for the most part with the requirements of the job position in terms of this dimension. Below we present the effect of *misalignment* in continuous dimensions (education, experience, region and log-wages) in a graphical way, in order to summarize the results of coefficients associated to the estimated polynomials in equation 3.

As noted above, for each dimension (except region) 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 (in the example, she has more years of experience than what is required in the job ad) and that the individual is underqualified if it's positive. For region, a categorical variable for each of Chile's 15 regions, we take the absolute difference between these variables, given the peculiar geography of Chile, where regions are organized consecutively, from north to south.[16]

In figure 4 we present graphically results of the effect of *misalignment*in levels of education, years of experience and region of residence on application decisions. The figure shows predicted application probabilities ($\hat{y}_{aw}$ from the estimates of equation 3), when a particular continuous dimension (z_{k_c}) 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 highly non-linear and depends on which value the other control variables take. The considered range for z_{k_c} is bounded by its sample mean plus and minus 1.5 times its standard deviation (only plus for the case of region).

As seen in figure 4, job seekers in both labor market states align themselves with some level of advertised requirements of job postings. This is represented by an inverted U shaped relationship between *misalignment* and application probability (all else constant) for education and experience, and by a decreasing line in the case of region.

[16]see https://en.wikipedia.org/wiki/Regions_of_Chile

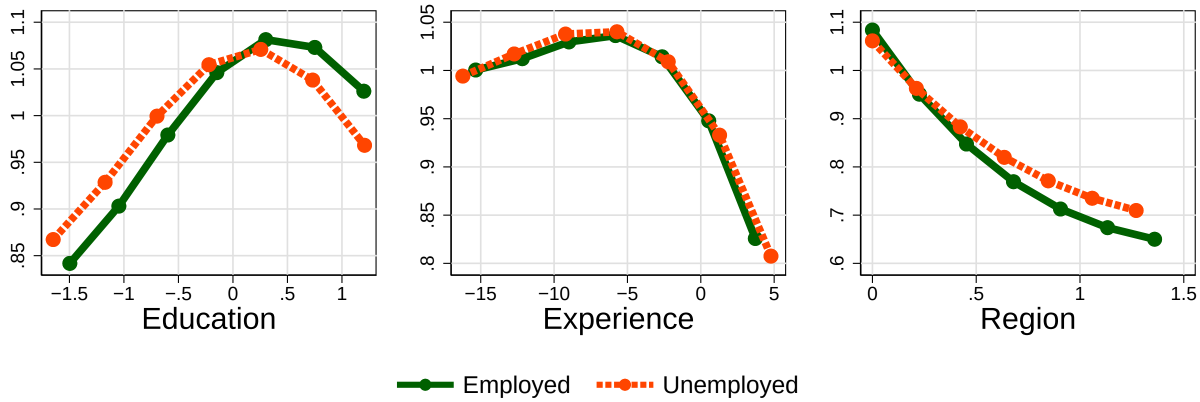

Figure 4: Predicted application probabilities, given results from eq. (3) and different levels of misalignment in the selected variable x (see main text for details). The rest of regressors are at the sample mean. Results are relative to unconditional application rates.

Both education and experience dimensions (left and center panels) are centered on a negative number (mid point in the x axis): the average *misalignment*in education is slightly below zero, while for experience, the number is slightly below negative five. This means that the average job seeker is slightly more educated and has five more years of experience than what is requested in the average job ad. In terms of region, the same inference cannot be made, since the horizontal axis reflects an absolute *misalignment*measure. Nevertheless, we can see that some applications in the sample are from individuals considering different regions but that considered regions are often not more than 1 region apart.

The alignment in both education and experience dimensions is not exact: for education, job seekers tend to maximize application probabilities at levels slightly above their own, while both employed and unemployed seekers apply with more frequency to jobs for which, on average, they have more experience than what is required by the job ad. The main reason for the average *misalignment* in this dimension is the fact our sampling strategy makes the average website user, someone attached to the labor force, with significant number of years of experience. For region, the figure shows clearly that the maximum application probability is centered at zero, meaning that most job seekers prefer job positions closer to them.

Given our exercise, we can analyze the presence of non-linearities for the education and experience dimensions. Figure 4 shows that job seekers react differently to ads for which they are under or overqualified. Moreover, this behavior is different in the case of education and experience. This follows from the inverted U shape curves, which are not symmetric around their peak. Job seekers tend to apply less to jobs to which they are over qualified in terms of education, compared to jobs for which they are under qualified: this is reflected on a steeper decline in application probabilities to the left of the peak than to the right of it, in the panel for education. On the other hand, individuals tend to apply less to jobs for which they are under qualified in terms of experience, as

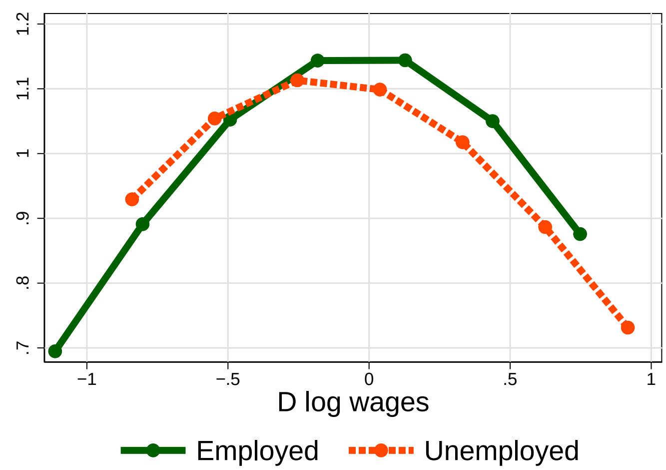

Figure 5: Predicted application probabilities, given results from eq. (3) and different levels of misalignment in log wages (expected to pay by the ad minus expected by worker). The rest of regressors are at the sample mean. Results are relative to unconditional application rates.

opposed to those jobs for which they are over-qualified (steeper slope to the left than to the right of the peak in the center panel of the figure). The figure also shows differences between preferences of unemployed versus the employed. From the education panel, we observe that unemployed seekers are less sensitive to being under/over qualified to job positions, which is reflected in the curve for them being flatter in the panel than that for the employed. As for experience, the differences between seekers in different labor states are minimal. For region, the drop in application probabilities is higher for employed seekers, meaning that unemployed seekers are relatively more likely to consider moving to obtain a job.

In figure 5 we show the same result, but for log-wages. Given that our estimates control for all other observables across job positions and job seekers, and that the regression controls for interactions, the *misalignment* in log-wages could be interpreted as *misalignment* in job and worker types: controlling for all observables, higher paying jobs and job seekers with higher earnings expectations must be of higher skill on average, and viceversa.

The figure summarizes well the overall search strategy that seems to be behind seekers in different labor market states. For the unemployed, the curve predicting applications reaches its peak slightly below zero. On the other hand, job seekers who already have a job direct their search towards job ads with types slightly above theirs, which is reflected in the solid curve in figure 5 reaching a peak at a positive level of the difference in log-wages. Thus, an intuitive takeaway from the evidence is that unemployed seekers try to maximize the probability of getting hired, while individuals performing on-the-job search, are interested in climbing the job ladder.

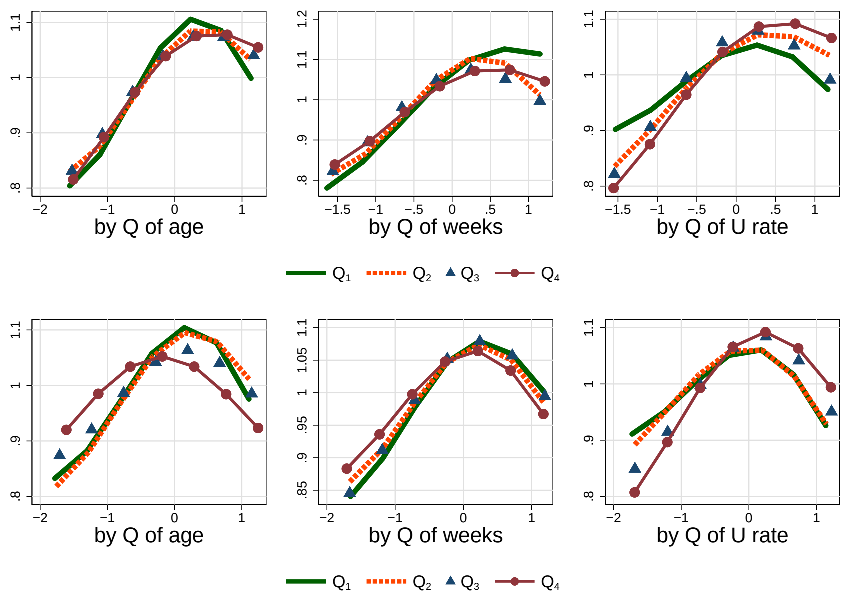

Figure 6: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from eq. (3) and different levels of misalignment in EDUCATION (see main text for details). The rest of regressors are at the sample mean. Results are relative to unconditional application rates. The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status and national unemployment rate.

Time and misalignment effects

In this section we ask how the effect *misalignment*in different dimensions at the time of searching for a job, interact with life-cycle, duration and business cycle effects. This could be seen as studying the interaction of the effects studied separately in the the two previous sections. Below, we separate estimation samples by quartiles of the three time variables: age of the worker, weeks in the current labor market status and level of aggregate unemployment at the time of the application decision. After estimating the regression in each of these sub-samples, we repeat the exercise of producing application probabilities by levels of *misalignment*. In the figures below, Q_1 to Q_4 represent the quartiles in ascending order, while the first row (group of three panels) shows effects for the employed, while the second row, for the unemployed.

In figure 6 we report the results of the exercise for the education dimension. In terms of age effects, the left two panels (labeled *by Q of age*) show that the effect of *misalignment*in education interacts heavily with age for the unemployed, but not much for the employed. The south-west panel shows that as age increases, unemployed seekers become less sensitive to educational *misalignment*with respect to job ads. Duration effects in the figure (center panels) are not significant, although slightly more pronounced for the employed sample.

The effect of business cycle effects is marked and similar for both employed and unemployed

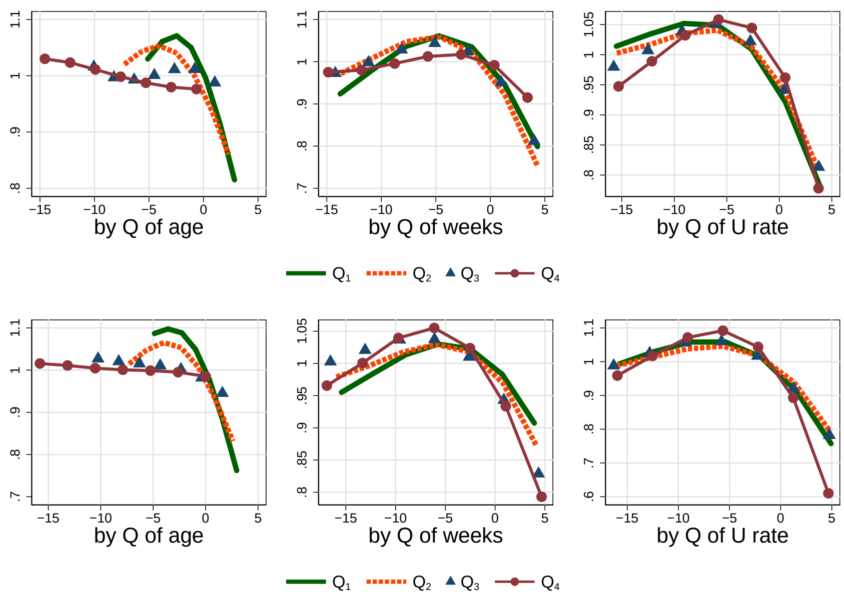

Figure 7: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from eq. (3) and different levels of misalignment in EXPERIENCE (see main text for details). The rest of regressors are at the sample mean. Results are relative to unconditional application rates. The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status and national unemployment rate.

seekers. When unemployment in the overall economy is higher (Q_4 versus Q_1) seekers in both labor market states concentrate more on the target education level, which is represented by a sliming of the predicted probability curves in the right panels of figure 6 for Q_4.

The case of *misalignment*in experience is displayed in figure 7. In the left two panels we observe that there are significant effects of age on how individuals align themselves with experience requirements of job ads: older workers exhibit increasing flatter curves in those panels, which implies that they don't align themselves to any particular level of required experience. A natural rationalization for this is that older job seekers have relatively more labor experience than the average job ad, so they do not find that information useful to make application decisions.

The effect of duration (two center panels) is different for employed (top panel) than for the unemployed (lower panel). For the earlier, our results show that longer tenure times have the effect of flattening out the application curve, so individuals performing on-the-job search become less sensitive to experience requirements the longer they have been employed in their current job. For the latter, the opposite occurs: unemployed seekers who have been longer in unemployment exhibit a slimmer application bell curve.

With respect to economy wide conditions, effects are also different depending on labor force status. For the employed (top right panel), higher unemployment rates are associated with comparatively lower application probabilities for the extreme cases of being over qualified in the experience

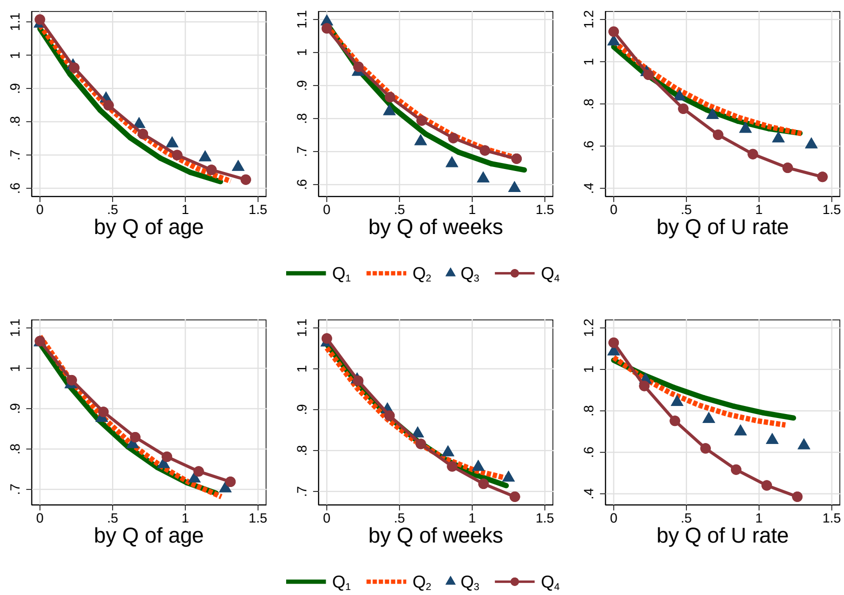

Figure 8: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from eq. (3) and different levels of misalignment in REGION (see main text for details). The rest of regressors are at the sample mean. Results are relative to unconditional application rates. The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status and national unemployment rate.

dimension: the left part of the application curve is lower under Q_4 than for the rest. For the unemployed (lower right panel). the effect is more marked for the extreme cases of under qualification of workers with respect to job ads.

In figure 8 we present the results for regional *misalignment*. The interaction between life-cycle and duration effects with *misalignment*in the location dimension are not salient, as reflected by the closeness of lines related to different quartiles in the left and center panels of the figure. However, there are significant effects with respect to aggregate conditions. In both top and lower right panels of the figure, we can observe that application probabilities fall significantly when the economy is experiencing high levels of unemployment, decreasing by almost a half. This is true for both the employed and the unemployed, but the effects are slightly stronger for the latter.

Finally, figure 9 shows results when we consider the (log) wage dimension. As stated before, one can view *misalignment*in this dimension as evidence of overall sorting of workers and job positions along types.

Results from the figure show small effects of either life-cycle, duration or business cycle conditions on application decisions, when individuals are considering *misalignment*in the log-wage dimension. Since we do find some effects above in terms of the other dimensions interacted with age, duration and unemployment rates, the lack of shifts in the application curves in figure 9 may reflect some averaging of all the effects, which in the end imply small total effects when individuals

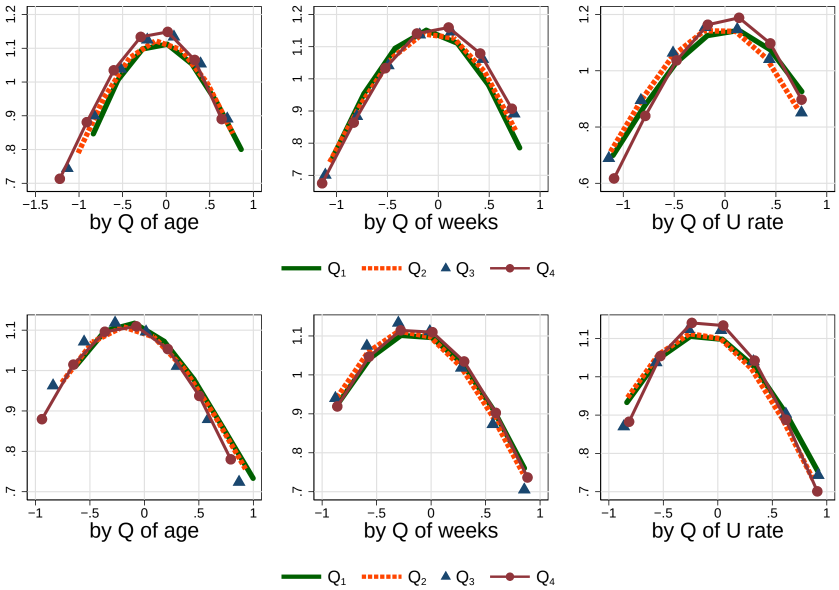

Figure 9: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from eq. (3) and different levels of misalignment in LOG-WAGES (see main text for details). The rest of regressors are at the sample mean. Results are relative to unconditional application rates. The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status and national unemployment rate.

consider *misalignment*in log-wages.

Conclusions

Using data from a Chilean job posting website, in this paper we uncover several facts regarding the nature of job search in an online setting. Given our unique setup, we can deconstruct behavior into two dimensions: an *extensive* dimension (number of applications sent) and an *intensive* one, where we analyze what affects the decision of sending an application at the margin.

For both dimensions, we find marked heterogeneity across those looking from unemployment and those performing on-the-job search, specially along the gender dimension. We find discouragement effects, in the sense that job search efforts decay with the time spent in the current labor force status (unemployment and employment). We also find new evidence regarding job search behavior during the business cycle: while the number of applications sent by individuals is clearly counter-cyclical (higher number when aggregate unemployment is down), aggregate conditions have a non-linear effect on the decisions of sending additional applications.

In the last section, we show how job seekers react to *misalignment* in key dimensions between own characteristics and characteristics required by job postings (level of education, years of experience, location, required occupation and log-wages) and find that there is significant alignment between requirements and characteristics. This alignment is affected by life-cycle, duration (time spent in the current labor force status) and aggregate business cycle effects.

References

Banfi, S., and B. Villena-Roldan (2016): "Do High-Wage Jobs Attract more Applicants? Directed Search Evidence from the Online Labor Market," Cea wp 327, University of Chile.

Banfi, S., and B. Villena-Roldán (2016): "Do High-Wage Jobs Attract more Applicants? Directed Search Evidence from the Online Labor Market," Discussion paper, Working Paper 327, Center for Applied Economics, University of Chile.

Carrillo-Tudela, C., and L. Visschers (2014): "Unemployment and Endogenous Reallocation over the Business Cycle," Cesifo working paper 4079.

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.
ΩŞahin, Song, Topa, and Violante

Şahin, A., J. Song, G. Topa, and G. L. Violante (2014): "Mismatch Unemployment," Forthcoming, American Economic Review.

Faberman, R. J., and M. Kudlyak (2013): "The Intensity of Job Search and Search Duration," mimeo, Federal Reserve Bank of Richmond.

Gee, L. K. (2015): "Information Effects on Job Application Rates by Gender in a Large Field Experiment," Mimeo, Tufts University.

Herz, B., and T. van Rens (2015): "Accounting for Mismatch Unemployment," Discussion paper, University of Warwick.

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): "Etude de la distribution florale dans une portion des Alpes et du Jura," 37, 547–579.

Jolivet, G., B. Jullien, and F. Postel-Vinay (2016): "Reputation and Prices on the e-Market: Evidence from a Major French Platform," *International Journal of Industrial Organization, forthcoming.*

Jolivet, G., and H. Turon (2014): "Consumer Search Costs and Preferences on the Internet," Mimeo, University of Bristol.

Kudlyak, M., D. Lkhagvasuren, and R. Sysuyev (2013): "Systematic Job Search: New Evidence from Individual Job Application Data," mimeo, Federal Reserve Bank of Richmond.

KUHN, P., AND H. MANSOUR (2014): "Is Internet Job Search still Ineffective?," *The Economic Journal*, 124(581), 1213–1233.

LEWIS, G. (2011): "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors," *The American Economic Review*, 101(4), 1535–1546.

MARINESCU, I. E., AND R. RATHELOT (2015): "Mismatch Unemployment and the Geography of Job Search," Mimeo.

MARINESCU, I. E., AND R. P. WOLTHOFF (2015): "Opening the Black Box of the Matching Function: The Power of Words," Discussion Paper 9071, IZA.

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 H. RUFFO (2015): "Optimal Life Cycle Unemployment Insurance," *American Economic Review*, 105(2), 816–59.

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