Unemployed Job Search Across People and Over Time: Evidence from Applied-For Jobs*

Jonas Fluchtmann[†] Anita M. Glenny [‡] Nikolaj Harmon [§] Jonas Maibom [¶]

Abstract

Using data on applied-for jobs for the universe of Danish UI recipients, we examine variation in job search behavior both across individuals and over time during unemployment spells. We find large differences in the level of applied-for wages across individuals but over time all individuals adjust wages downward in the same way. The decline in applied-for wages over time is descriptively small but economically important in standard models of job search. We find similar results when examining variation in the non-wage characteristics of applied-for jobs and in the search methods used to find them. We discuss implications for theory.

Keywords: job search, unemployment

JEL: E24, J64

*This paper has benefited from numerous questions and comments from participants at several seminars, conferences and workshops. We thank Alexander Koch, Rune Vejlin, Andrea Weber, Roland Rathelot and Daniel le Maire for comments on early versions of the paper. We are grateful to the Danish Agency for Labor Market and Recruitment and the ECONAU project database (Department of Economics and Business Economics, Aarhus University) for making their data available. This paper was supported by the Independent Research Fund Denmark grant "How do unemployed job seekers apply for jobs and can we help them do better?".

[†]Aarhus University; present: OECD; E-mail: jonas.fluchtmann@oecd.org; The views expressed here are those of the authors and do not necessarily reflect the official views of the OECD or its member countries.

[‡]Aarhus University; E-mail: glenny@econ.au.dk;

[§]University of Copenhagen ; E-mail: nikolaj.harmon@econ.ku.dk

[¶]Aarhus University and DTMC; E-mail: maibom@econ.au.dk;

1 Introduction

When unemployed workers engage in job search, they make important choices about which jobs to target and how to find them. In this paper we ask how these choices vary both across workers and over time. Do some workers make very different search decisions than others? Does a given worker change behavior after having searched unsuccessfully for some time? Past work dealing with these questions has been constrained by data availability. While administrative data often provide comprehensive information on job finding, these types of data typically have no direct information on job search behavior. Meanwhile, existing micro-data on search behavior typically only cover a small and selected sample of individuals or search behavior, such as the users and job applications made on a single online platform.

In this paper, we overcome these data constraints using novel administrative data on applied-for jobs. Since 2015 all Danish UI recipients have been required to document applied-for jobs electronically with the Danish employment agency. By linking these data with administrative data on UI recipients, we create a panel data set which contains information on applied-for jobs during each month of the UI recipients' unemployment spells. A key feature of these data is their coverage. By construction, the administrative data include the universe of UI recipients. In addition, a particular advantage of the Joblog application data is that they cover all types of applied-for jobs rather than being limited to job applications made via a certain channel or platform, or being limited to a certain subset of potential jobs. As we expand on in the paper, we estimate that the Joblog data cover between 69 and 80 percent of UI recipients' applied-for jobs, and that the covered jobs are highly representative. We use these data to examine differences in UI recipients' search behavior across individuals and over time. In addition to focusing on the wages of applied-for jobs, we also examine hours, geographical proximity, industry and occupation. Because job seekers provide information about how they found each registered job, we also examine the methods which job seekers use to search for jobs.

We structure our analysis around predictions from a benchmark random job search model. In the model, unemployed workers are heterogeneous and apply for jobs in a stationary environment with search frictions. The distribution of jobs is fixed, and jobs are characterized by the wages they pay. The model generates three key predictions that we take to data. First, persistent worker heterogeneity generates permanent differences in applied-for wages across individuals. Second, individuals never change behavior over time, but applied-for wages still vary with the completed unemployment duration because the pool of unemployed changes over time - so-called dynamic selection. Third, outside of wages, the model generates no systematic variation in search methods or applied-for job characteristics.

We begin our empirical analysis by looking at applied-for wages. We find large variation in applied-for wages, most of which reflect persistent differences across individuals. The standard deviation in average applied-for wages is 13 percent and the month-to-month correlation is above 0.7. This finding is in line with the prediction of the benchmark model with substantial time-invariant worker heterogeneity.

Next, we examine variation in job search behavior over time. To address dynamic selection, we exploit the longitudinal nature of our data and estimate a two-way fixed effects model with individual- and time-fixed effects. This model turns out to fit the data well, also compared to more flexible specifications with differential time effects. The estimated time effects show that individuals continuously lower their applied-for wages over time. The implication is that while there are large level differences in applied-for wages across individuals, over time everyone adjusts their applied-for wages downwards in the same way. Over the first year of unemployment, applied-for wages decline by 1 percent in total. This contradicts the benchmark model in which individuals do not change behavior over time. The use of longitudinal data to address dynamic selection turns out to be crucial for establishing this result. In the raw data, most of the decline in applied-for wages over time is obscured by dynamic selection.

Examining other dimensions of search behavior, we find patterns similar to those for applied-for wages. Across individuals, there are large and persistent differences in both applied-for hours, occupation, industry, geography and search methods. Two-way fixed effects models again fit these data well suggesting large level differences across individuals but similar adjustment patterns over time. Mirroring the results for wages, the time patterns show that workers target less and less attractive jobs over time also in these other dimensions. For search methods, workers gradually rely more and more on formally posted vacancies over time, at the expense of informal search channels such as social networks. The benchmark model does not generate systematic differences along these dimensions because it models search decisions as only being related to wages.

We finish our empirical analysis by examining the correlation between applied-for wages, total realized unemployment duration and other observables. We find a U-shaped relationship between applied-for wages and total realized unemployment duration: Throughout the spell, both individuals who exit UI very quickly and individuals who remain long-term unemployed apply for systematically higher wages than individuals who exit unemployment in the medium term. This pattern is rationalized by the benchmark model if individuals who exit unemployment fast face higher job offer arrival rates, while the long-term unemployed are characterized by being very patient or having a high value of unemployment. Observables are capable of explaining 36 percent of the variation in average applied-for wages and most of the relationship between applied-for wages and total spell length.

In the final part of our paper, we discuss the theoretical implications of our results. We first examine the quantitative importance of declining applied-for wages over time. Descriptively, the observed 1 percent-per-year decline is unimportant relative to the standard deviation of applied-for wages of 13 percent. To assess the economic importance of the decline in applied-for wages, we examine how the observed decline may impact job finding. We derive a simple formula for the elasticity of the job finding hazard with respect to the average applied-for wage which holds in a wide range of commonly used search models. Implementing this formula on our data, we find that declines in applied-for wages have important consequences for job finding. If individuals did not change their applied-for wage while unemployed, job finding rates a year into unemployment would be 4-8 percent lower than what we see in the data.

Second, we discuss ways to rationalize the time patterns in search behavior that we find, also across non-wage job characteristics and search channels. Non-stationary job search models with negative duration dependence or UI benefit exhaustion predict that workers should be willing to accept less attractive jobs later in their unemployment spell. This can explain our finding that individuals lower their applied-for wages over time. If the models are extended so that workers consider not only wages but also non-wage job characteristics, it can also explain why we find systematic adjustments in applied-for non-wage job characteristics. Alternatively, stock-flow matching and other models in which workers gradually exhaust their search prospects generate similar time patterns in behavior and can also explain why workers shift to rely more on formally posted vacancies as other search channels are exhausted. Finally, learning models or models with reference-dependent preferences also predict systematic changes in behavior over time that could fit our results.

Finally, the finding that all dimensions of job search behavior are well described by models with individual- and time-fixed effects also has implications for both theory and empirical work. To match our results, search models must include worker heterogeneity and dynamics in a way that generates large dispersion in search behavior across workers but parallel adjustments over time. We highlight one common modeling approach that satisfies this restriction as an example. For empirical work, the very stable parallel trends across groups serve as a validation of the parallel trends assumption that underlie many difference-in-difference designs employed in the empirical literature on job search.

The existing work most closely related to ours is Marinescu and Skandalis (2021). Using administrative data from France linked with information on applied-for jobs from the job search platform of the French Public Employment Services, they also study changes in job search behavior over time but focus their analysis on changes in behavior around benefit exhaustion. We complement this work by focusing both on heterogeneity across individuals and over time, and by exploring a much wider set of applied-for job characteristics.

A number of seminal papers use survey data or data from online platforms to document patterns of job search behavior. Highly related to our results is Kudlyak et al. (2013). Using data from an online platform, they also provide evidence that workers over time apply to jobs that pay lower and lower wages, as proxied by the education level of the other applications. Other examples are DellaVigna et al. (2021); Krueger and Mueller (2016, 2011); Mueller et al. (2021) using survey data and Marinescu and Rathelot (2018); Banfi and Villena-Roldán (2019); Banfi et al. (2019a,b); Kuhn and Shen (2013); Faberman and Kudlyak (2019) using data from online platforms. Relative to these papers, a main strength of our analysis is the coverage of the data we use. Our data are not limited to covering search activities on a single platform and do not suffer from concerns about non-random attrition.

Other recent, related papers include Le Barbanchon et al. (2020) and Fluchtmann et al. (2019) who focus on gender differences in job search, Belot et al. (2019), Altmann et al. (2018), Skandalis (2018) and Gee (2019), who study the impact of information about e.g. job opportunities on job search, and Lichter and Schiprowski (2021), Le Barbanchon et al. (2019) who study the effect of changes in the UI benefit system.

The rest of the paper is structured as follows: In Section 2 we present the benchmark search model that we use to guide our analysis. In Section 3 we present the institutional setting and the data sources used in the paper. This includes an extensive discussion of the validity and content of our data on applied-for jobs. In Section 4 we present our empirical analysis and results. In Section 5 we draw implications of the results for theory. Section 6 concludes.

2 Motivating theory

We use a benchmark model of unemployed job search to motivate and guide our empirical analysis (see e.g. Mortensen, 1977; McCall, 1970; Rogerson et al., 2005). Workers entering unemployment search for jobs in continuous time. From the worker's perspective, available jobs differ only in their wage offers $w \geq 0$ which follow some distribution F. Unemployed workers receive a flow utility of b. At Poisson rate λ, unemployed workers receive a job offer drawn randomly from F and then decide to accept or reject. If a worker accepts a job offer, they are assumed to stay employed forever. Workers are forward-looking, discount the future at rate r and expect all parameters to remain unchanged.

Workers' decisions about search behavior are captured entirely in the job acceptance decision. This in turn can be summarized by the choice of a reservation wage, w^R, such that all offers $w \geq w^R$ are accepted. Letting U denote the value of being unemployed, optimal behavior implies setting the reservation wage equal to the flow value of unemployment:

$$w^R = rU \tag{1}$$

The value of unemployment solves the following continuous time Bellman equation (see also Rogerson et al. (2005) equation 11 and the surrounding discussion):

$$rU = b + \lambda \int_0^\infty \max\left(\frac{w}{r} - U, 0\right) dF(w) \tag{2}$$

Finally, the hazard rate at which workers find jobs and leave unemployment is:

$$h = \lambda\left(1 - F(w^R)\right) \tag{3}$$

The rate at which workers leave unemployment thus depends on the job offer arrival rate, λ, and the probability of getting a job offer which is above the reservation wage, $1 - F(w^R)$.

2.1 Search behavior, reservation wages and job applications

The benchmark model equates search behavior simply with the choice of a reservation wage. Our empirical analysis relies on data about job applications, not reservation wages. We view this mostly as a strength. Whereas job applications are direct measures of search behavior, the combination of a reservation wage and exogenous wage offer arrivals is typically viewed primarily as a useful theoretical abstraction.[1] Linking our job application data to the benchmark model, however, does require us to take a stance on how applications relate to search behavior in the model.

Our preferred approach is to assume that job applications made by a worker represent random draws from the part of the job offer distribution that lies above the worker's reservation wage. Letting a index applied-for jobs observed in the data and letting w_a be the wage of applied-for job a, we assume

$$w_a \sim F(w|w \geq w^R) \tag{4}$$

In Appendix A.2, we show that this assumption follows from a standard reinterpretation of the benchmark model that explicitly considers job applications.[2]

Let $w^* \equiv E(w_a)$ denote the average wage among applied-for jobs. This will be a key variable in our empirical analysis. From 4 it follows that w^* corresponds to the expected wage among acceptable job offers:

$$w^* = E\left[w|w > w^R\right] = \tfrac{1}{1-F(w^R)} \int_{w^R}^{\infty} w \, dF(w) \tag{5}$$

As long as w^R is in the support of F, Equation 5 implies that there is a monotonic, one-to-one mapping between the average applied-for wage and the reservation wage. Qualitative results for applied-for wages thus translate directly to the reservation wage and vice versa. In what follows, we therefore use the term *targeting high wage jobs* as a short-hand for having a higher reservation wage and - equivalently - higher applied-for wages.

Appendix 6 provides additional discussion of the link between job applications and the benchmark, including an alternative interpretation of Equation 5 as stemming from an equivalent model of directed search (see e.g. Nekoei and Weber (2017); Marinescu and Skandalis (2021)).

[1] The reservation wage does not correspond to any directly observable behavior but pertains to a counterfactual about which wage offers a job seeker would accept. Empirical measures of reservation are typically elicited through surveys with various measurement issues (see e.g. Cox and Oaxaca (1990); Kesternich et al. (2022)). For example, Krueger and Mueller (2016) report that a fair share of job seekers accept (reject) job offers below (above) their stated reservation wage. A further challenge is that the reservation wage describes the minimum acceptable wage offer for a "reference job", however the content of the "reference job" may vary across respondents causing problems when other characteristics besides the wage affect the total utility of the job (see for example the discussion in Hall and Mueller (2018)). We discuss the role of non-wage job characteristics further in Section 5.3.

[2] Instead of a fixed arrival rate of job offers, we assume that the worker learns of a potential job opportunity at some fixed rate, including learning the potential job's per-period wage w, drawn randomly from F as before. Upon learning about the job, the worker now faces a decision about whether to apply for the job or not. This decision is summarized in an (application) reservation wage w^R. If they apply for the job, there is a fixed probability p that the application is successful and they get the job, otherwise they stay unemployed. See Appendix A.2.

2.2 Worker heterogeneity and variation in search behavior across individuals

In the benchmark model, differences in search behavior can arise from heterogeneity in the worker parameters b, λ and r.[3] These differences follow standard comparative static results on reservation wages.

If differences in UI benefits or other factors cause some individuals to have a higher flow value of unemployment, b, these individuals will have a higher value of unemployment and will thus have a higher reservation wage and a lower hazard rate out of unemployment. A similar result holds for patient individuals with a lower discount rate r. Variation in b or r across individuals thus implies that we should see some individuals target higher-paying jobs while simultaneously exiting unemployment more slowly.

If some individuals are more effective at searching or at securing job offers from employers, these individuals may have a higher arrival rate of job offers, λ. This again raises their value of unemployment and their reservation wage. Under standard regularity conditions (see e.g. Rogerson et al. (2005) footnote 9), however, their higher arrival rate of offers will still imply a higher hazard rate. Variation in λ thus allows for the possibility that some individuals target higher paying jobs while simultaneously exiting unemployment faster.

In our empirical analysis we use data on applied-for wages to examine variation in the wages workers target. Using linked data on job finding and rich worker characteristics, we also further relate this variation to both job finding rates and observable characteristics..

2.3 Dynamic selection and variation in search behavior over time

Because of the stationary environment, the benchmark model does not generate variation in individual job search behavior over time: i.e. a given worker does not target different jobs over time. With heterogeneity across individuals, however, the model *does* predict that job search behavior differs across people with different completed unemployment duration. This is due to the well-known issue of dynamic selection: As discussed in the previous section, variation in b or r implies that some workers target higher-wage jobs than others while simultaneously remaining unemployed longer on average. Because such individuals will make up a larger share of the long-term unemployed, the group of long-term unemployed will target higher wages on average. The reverse will be true if variation in λ causes some individuals to target higher wages while exiting unemployment faster.

In our empirical analysis, we use longitudinal data on applied-for jobs to follow search behavior for the same individual over time. This allows us to address dynamic selection and identify whether individual search behavior changes over time. Importantly, because our data cover the universe of UI recipients for the full duration of their UI spell, we can do this without facing issues of survey attrition over time.

[3] Mechanically, one could also include worker heterogeneity in the distribution that wage offers are drawn from F. This, however, amounts to assuming that workers differ in the types of job offers they attract and thus breaks with the standard interpretation of random search in which F represents the objective distribution of vacant jobs that workers search for completely at random. In Appendix A.2 we discuss alternative interpretations of our data and benchmark model, building on the idea that job search is directed instead of random.

2.4 Other dimensions of search behavior

The benchmark model focuses on job acceptance decisions and assumes that wage level is the only utility relevant difference across jobs. This implies that search decisions deal exclusively with (reservation) wages. Our data contain rich information about other dimensions of search, however. We use these to test for broader variation in search behavior.

In terms of job characteristics, we focus on two dimensions. First, we examine whether there is systematic variation in the hours of jobs that individuals target. Second, we examine systematic variation in what we refer to as the scope of search - whether individuals target jobs that are closely related to their past industry or occupation, and whether jobs are geographically close to where they live.

Finally, our data also contain information on the method by which job seekers found the jobs they are applying for. The benchmark model is silent on the method by which workers search, however, richer theories can generate variation in this dimension. In our analysis, we thus also examine systematic differences in the search methods that workers use.

3 Data and institutional setting

Over the next subsections, we first present the details of our empirical setting and the data set we use which combines data on applied-for jobs (the so-called Joblog data) with other administrative data sources (see also Fluchtmann et al. (2019), in which we use the same data sources to study gender differences in job search).

3.1 The Danish UI system and the Joblog application data

The Danish UI system is based on voluntary membership. To be eligible for UI, workers are required to sign up and contribute to one of the 24 UI funds sufficiently well in advance of becoming unemployed. The vast majority of Danish workers satisfy these eligibility requirements.[4] When a UI eligible worker becomes unemployed, benefits are available for up to two years. UI benefits are determined at a replacement rate of 90 percent of previous income and a cap of 18.500 DKK (2.500 euro in 2017) per month. The cap is binding for the majority of workers. For workers who exit unemployment, the two-year eligibility period resets after one year of employment. For workers who run out of UI benefits, social assistance benefits are instead available. Social assistance benefits are means tested and thus not available for individuals with e.g. savings or other assets.

To remain eligible for UI while unemployed, UI recipients have to be actively searching and applying for jobs. Further, they regularly have to document that they are satisfying eligibility criteria, i.e., show examples of submitted applications etc. Since 2015 this documentation has been

[4]In 2015, 76 percent of Danish employees were members of a UI fund while over 70 percent of the gross unemployed were UI recipients. Among the residual group of gross unemployed, more than 20 percent receive means-tested social assistance (which amounts to around 60-80 % of the maximum UI benefits) and are therefore likely to have exhausted UI prior to this (see e.g. Danish Economic Council, 2014).

centralized through an online system called *Joblog*. Through the Joblog platform UI recipients fill in a form describing the particular job they have applied for. It is mandatory to provide information on the applied-for job, including job title and hours (part-/full-time) and about the potential employer, including firm name and address. In addition, the job seekers must also provide information on how they found and applied for the job. As we describe further below, it is the information entered on these forms that we use to measure job search behavior and link with additional administrative data.

As a general rule of thumb, UI recipients are instructed that they need to register somewhere between 1.5 and 2 applications per week in the Joblog system to maintain eligibility. Failure to comply with these documentation requirements would ultimately result in sanctions in the form of lost or reduced UI payments.[5] UI recipients thus face a clear economic incentive to comply with the requirements and register submitted job applications in Joblog. As we discuss further below, these incentives have resulted in a very high level of usage and correspondingly a high level of coverage for our data.

3.2 Selecting the analysis sample

Our baseline sample is constructed from administrative data on UI payments and consists of all individuals of Danish origin who start unemployment spells between September 2015 and September 2017, and who are eligible for the full two years of UI benefits. The latter requirements ensure that all in our analysis face the same incentives in regards to the timing of future benefit exhaustion.

We follow each new UI recipient for a year or until the UI spell ends, which we define as having 4 consecutive weeks where no UI is paid out. This definition implies that a UI exit can be associated with entering employment or other public benefits, e.g. sickness benefits. In practice, however, the vast majority of UI spells end with employment. Since our data end in September 2017, we note that there are a limited number of spells that are right-censored (see Table 4 in Appendix A.1). None of our results hinges on the inclusion of these censored spells.

By focusing on the first year of unemployment, our analysis differs from much of the existing literature on the dynamics of job search that emphasize behavior later in the unemployment spell, around the time UI benefits expire. Our main motivation for focusing on the early part of an unemployment spell has to do with quantitative importance. While studying behavior at the time of benefit exhaustion can be extremely useful for testing different theories of job search (see e.g. Card et al., 2007; Ganong and Noel, 2019; DellaVigna et al., 2021; Marinescu and Skandalis, 2021), benefit exhaustion is in fact a relatively rare event in many settings - particularly in settings with relatively generous UI duration (Danish Economic Council (2014)). For example, in the Danish

[5] In the case of non-compliance with the job search requirements, UI recipients will typically be given a short time period to prove eligibility and register previously unregistered (or ongoing) job search after which the UI fund will make its final assessment. The size of the sanctions ranges from a loss of benefits for a couple of days to a permanent loss of benefits depending on the severity of the non-compliance. In cases where registered job applications are not considered adequate (due to e.g. an assessed risk of proforma search or fake applications), similar requirements apply. See also Online Appendix B.1 for additional institutional details.

setting more than 80 percent of individuals exit the UI system within the first year.

For each UI recipient and unemployment spell, we use a unique person identifier to identify all applied-for jobs that they have registered in the Joblog system during their unemployment spell. We further use this person identifier to include a wide range of other administrative data sets maintained by Statistics Denmark. These data sets include information on demographics, education, public benefit payments and employment, including information on occupation, wages and firm identifiers for the employing firms.

To arrive at our analysis sample, we impose four additional restrictions. First, we impose the restriction that the UI spell lasts at least 8 weeks. We do this to remove individuals who are de facto making a job-to-job transition but who are temporarily receiving UI while waiting for their new job to start. Second, we discard job applications in the last 4 weeks of a spell. This gets rid of applications that UI recipients are making after successfully landing a job, but before this job has actually started.[6] Third, we require that the individual has registered at least one applied-for job during the observed unemployment spell so that some information on search behavior is available. Fourth, after imposing our other restrictions, only 434 individuals in our data show up with multiple UI spells. We keep only the first spell for these individuals. In Online Appendix B.2 we provide additional discussion of our sample construction and data.

Our final analysis sample consists of 127,695 individuals and over 4 million applied-for jobs. In Table 4 in Appendix A.1 we provide some summary statistics on our final sample for the analysis. We report statistics for both the full sample and across different unemployment duration groups.

3.3 Measuring job search behavior

We use the information on applied-for jobs in Joblog to construct monthly measures of job search behavior. For each applied-for job in the data, the UI recipient reports a range of information directly. Additionally, we link the applied-for job to a firm identifier in the administrative data based on the reported firm name and address, and to an occupation based on the reported job title. In the source data, we successfully match 86 percent of applications to a firm and 82 percent to an occupation (see Online Appendix B.3 for additional details).

The first measure of search behavior we focus on is the wages of applied-for jobs. The data do not contain direct measures of the wage an applicant would have been paid in each applied-for job. Instead we use realized data on wage payments for new hires to predict the (log) wage for each potential job based on the observable characteristics of the job, including the occupation of the job and the firm effect of the employer from an AKM wage model (Abowd et al. (1999), see Online Appendix B.4 for details). For each UI recipient and each month, we use these predictions to compute the average wage of the applied-for jobs in our data. As a robustness check we also report results where we directly use the firm fixed effect to capture the wage level at the applied-for

[6] As consequence of the wording of the Danish UI rules during our sample frame such individuals were in principle required to both apply for and register applications in Joblog, despite the fact that they had a new job lined up with a known start date. They obviously face a very peculiar set of incentives in their application decisions. By dropping applied-for jobs in the last 4 weeks before a new job start we get rid of these applications (See Online Appendix B.2).

firm (see Online Appendix C.1).

In addition to wages, we also examine the hours of the applied-for job. In particular, we measure the share of applied-for jobs that are reported as being full-time by the UI recipient. We also examine the scope (or broadness) of search in terms of geography, industry and occupation. Based on driving time data from Harmon (2013), we measure the commute time to each applied-for job based on the UI recipients' zip codes and the location of the job. For occupation and industry, we determine whether each applied-for job is in an occupation and/or industry that is related to that of the UI recipient's previous job. We measure relatedness based on the O*Net Related Occupation Matrix and an industry transition matrix for the labor market. For our main analysis, we collapse these into a single monthly 'scope of search' index based on the average commute of applied-for jobs, the share of applied-for jobs that are in a related occupation, and the share of applied-for jobs that are in a related industry (see Online Appendix B.5 and C.3.1 for additional details on variable construction and for robustness checks using alternative measures of related occupations).

Finally, we examine the search channel used to find each applied-for job, focusing specifically on the distinction between formally posted vacancies and more informal channels. In the raw Joblog data, job seekers directly report whether they found the job via a publicly posted vacancy, whether they heard about the job through their social network, whether they were initially contacted by the firm (or other headhunters), or whether they simply applied to a firm of interest without any indication that the firm had open positions (what we refer to as a 'cold contact'). In our main analysis, we examine the share of applied-for jobs found via a publicly posted vacancy each month (we show results for the other outcomes in Online Appendix C.4).

3.4 Coverage and validity of the job search data

Relative to many other data sets with information on job search, a key advantage of the data sources we use is their coverage. By definition, the administrative data include information on the universe of UI recipients in Denmark. In addition, a particular advantage of the Joblog application data is that they cover all types of applied-for jobs rather than being limited to job applications made via a certain channel or platform, or being limited to a certain subset of potential jobs. Since UI recipients themselves are responsible for registering the applied-for jobs, the coverage of these data still warrants attention however.

As noted previously, a key feature of the institutional setting is that UI recipients face clear economic incentives to comply with recording requirements and register applied-for jobs. Indeed the data suggest strong compliance with the requirement to register 1.5-2 jobs per week. In the raw data, almost all UI recipients register jobs. Moreover, the average number of applied-for jobs per week is just above 1.5 (see Table 4), and the distribution of applications per week has a clear mode of around 2. In Online Appendix B.6 we use auxiliary survey data to examine which share of all applications is registered in the Joblog system. We find a high level of coverage: Survey results suggest that UI recipients register between 69 and 80 percent of all their applied-for jobs.

Another reassuring institutional feature is that UI recipients face no formal incentives to selectively register some applications over others in Joblog. This suggests that the applied-for jobs included in our data are likely to be a representative subset of all applied-for jobs. In Online Appendix B.6 we subject the data to a range of validity checks to verify this. These checks exploit the fact that - independently of the application data - we also observe actual job outcomes. Reassuringly, the checks all suggest that our data are highly representative and that registered applications are highly predictive of later job outcomes.

The institutional setting around our data does give rise to two important caveats however. First, the strong compliance with reporting requirements makes the number of applied-for jobs uninformative about actual search behavior; individuals who apply for more jobs than the reporting requirements face no formal incentive to register these additional jobs. Indeed both the raw data and supplementary surveys suggest that they often do not. In our analysis, we therefore make no attempt to infer individual search effort from the observed number of applied-for jobs as has been done in other settings (Faberman and Kudlyak (2019); Marinescu and Skandalis (2021); DellaVigna et al. (2021))[7]

The second potential caveat has to do with external validity. If the stark search requirements cause Danish UI recipients to apply to more jobs or to different types of jobs, our results may not generalize to other settings. Two facts alleviate this concern somewhat however. First, although the electronic reporting is quite unique, the search requirements in the Danish context are in fact not so different from what is seen in many other countries. For comparison, similar types of activity requirements exist in several other UI systems, see for example Marinescu (2017) and McVicar (2020).[8] Second, the requirement to apply for and document 1.5-2 jobs per week appears fairly similar to the number of jobs unemployed job seekers apply to also in other contexts. For example, our numbers are similar to Faberman et al. (2021) who report that unemployed job seekers send slightly more than 2 applications per week on average in the US.

4 Empirical Results

We begin our analysis by examining the overall variation in our measures of job search behavior in Table 1. The top left of the table shows the mean and standard deviation of the monthly average applied-for wage in our analysis sample. The raw variation in search behavior is substantial; the standard deviation of applied-for wages is 13 percent (log points).

Viewed through the lens of the benchmark search model from Section 2, this variation is consistent with substantial worker heterogeneity in model parameters such as the discount rate, the flow value of unemployment and/or the offer arrival rate. Of course, an alternative explanation is that

[7] For completeness Figures 3 and 4 in Online Appendix B.2 show distribution plots of the number of weekly and monthly applications as well as the evolution in the monthly number of submitted job logs over time in unemployment. See also the discussion of the results in Online Appendix B.2.

[8] For example Marinescu (2017) reports that: "*A standard level of monitoring may require the unemployed to contact about two employers per week, and report this to the unemployment agency, as in the US state of Maryland [2], or in Switzerland [3].*" Therefore the Danish requirements are in fact quite similar to requirements in other settings, although the enforcement may of course be perceived as stricter.

most of the variation here is idiosyncratic, stemming from randomness in the search process and/or various forms of measurement error. As a simple check on this, we compute the month-to-month correlation in the average applied-for wage across individuals in our data. If the variation in applied-for-wages reflects mostly random search or idiosyncratic measurement error, this correlation should be small. Conversely if the variation is driven by permanent worker heterogeneity, the correlation should be high.[9] As shown in the upper-right corner of Table 1, the month-to-month correlation in the average applied-for wage is in fact high, 0.77.

The rest of Table 1 shows corresponding results for our other dimensions of search behavior. Across all measures, we see substantial variation. Moreover, the month-to-month correlations show that this variation can not simply be explained by idiosyncratic factors. This systematic variation in other dimensions of search cannot directly be explained by the benchmark model, which considers only wage differences across jobs.

4.1 Variation in applied-for wages over time

Next we examine variation in applied-for wages over the unemployment spell. We start by examining the raw variation in behavior across unemployment durations. A convenient way to do this is to estimate a simple time fixed effects regression. Letting i index spells, t index months of the unemployment spell, and letting $\bar{w}_{it}$ be the average applied-for-wage in month t of the ith spell, we estimate the following regression on our analysis sample

$$\log \bar{w}_{it} = \tau_t + \varepsilon_{it} \tag{6}$$

We refer to the set of time fixed effects, $\tau_1, \tau_2, ..., \tau_{12}$, as duration effects. Of course, with no other regressors, the estimate for the t-th duration effect will simply equal the raw mean of $\bar{w}_{it}$ for month t (across all spells that last at least until this month).

The blue line in Panel A of Figure 1 shows the estimated duration effects from Equation 6. Over the first four months, there is a decrease in the average applied-for wage of around 0.5 percent (log points), but from month four and onward the average applied-for wage is more or less flat.

As discussed in Section 2.3, estimates from Equation 6 may be subject to dynamic selection that obscures individual changes in behavior over time. If some individuals permanently target higher-wage jobs but exit unemployment slower or faster than others, these individuals will make up a larger or smaller share of the sample over time. We use the longitudinal nature of our data to address dynamic selection. Specifically, we expand Equation 6 with an individual fixed effect:

$$\log \bar{w_{it}} = \alpha_i + \tau_t + \varepsilon_{it} \tag{7}$$

[9]Formally, assume that the log average applied-for wage, $\bar{w}_{it}$, is the sum of a individual specific component, α_i, and (possibly) an idiosyncratic component, ε_{it}, that is i.i.d. The month-to-month correlation in $\bar{w}_{it}$ then shows which share of the variation in applied-for wages is explained by individual heterogeneity:

$$\frac{Cov(\bar{w}_{it}, \bar{w}_{it+1})}{Var(\bar{w}_{it})} = \frac{Var(\alpha_i)}{Var(\alpha_i) + Var(\varepsilon_{it})}$$

Table 1: Summary stats on applied-for job

	Full sample	Month-Month
	Mean	Correlation
Avg. applied-for wage (log)	5.18	0.77
	(0.13)	
Avg. commute time	44.72	0.69
	(35.03)	
Share full-time	0.89	0.64
	(0.22)	
Share occupation unrelated to prev. job	0.43	0.77
	(0.40)	
Share industry unrelated to prev. job	0.52	0.63
	(0.38)	
Scope of search index	0.08	0.75
	(0.999)	
Share publicly posted vacancy	0.74	0.66
	(0.33)	
Share found via network	0.06	0.43
	(0.16)	
Share cold contacts	0.19	0.64
	(0.29)	
Observations	575,504	355,177

Note: This table reports the mean, standard deviation (in parenthesis), and month-to-month correlation for different monthly measures of job search behavior. Observations are individuals-by-month. See Section 3.3 for further details on the different job characteristics. The number of observations reported in the table is for average applied for wages (note that for some months we do not have data for all measures, see Table 7 for the number of observations for each measure.

In Equation 6 individuals have different initial levels of applied-for wages as captured by α_i, but then all adjust their applied-for wages in the same way over time as captured by τ_t. By absorbing permanent differences in the level of applied-for wages, the individual fixed effect α_i addresses dynamic selection on levels; i.e., the estimated duration effects will not be affected if individuals with permanently higher applied-for-wages exit unemployment faster or slower (we return to the possibility of dynamic selection on changes in Section 4.3).

When estimating Equation 7, we normalize τ_1 to equal the sample mean of the dependent variable in month 1. This implies that the estimated duration effect for month 1 will be numerically the same as in the regression without individual fixed effects (Equation 6).[10]

The red line in Panel A of Figure 1 shows estimated duration effects based on Equation 7. We

[10] Because Equation 7 includes both a individual fixed effect and a full set of duration effects, it is necessary to impose a normalization for estimation. Throughout most of our analysis, we normalize τ_1 to equal the mean of the outcome variable in the first month of unemployment. This makes it straightforward to compare the estimated changes over time to the level of the outcome variable in the overall sample. Of course, the choice of normalization has no effect on the estimated changes in behavior. Normalizing τ_1 to zero (as is more common) would simply shift down all our estimated coefficients by the sample mean of the outcome in the first period.

see a different pattern than before. Instead of flat-lining after four months, applied-for wages now continue to decrease across all months. Over the first year of unemployment, the average applied-for wage decreases by about 1 percent. This systematic change in search behavior throughout the unemployment spell contradicts the predictions of the benchmark model where individuals never change their search behavior over time. In Online Appendix C.1, we show that the time profile in Figure 1 is similar if we replace the average applied-for wage with the average AKM firm fixed effect of the applied-for firms or even alternative measures of a firm's "type". The change in applied-for wages documented in Figure 1 thus in part reflects a change in the type of firms workers apply to.

We postpone our discussion about the quantitative economic implications of the observed drop in applied-for wages until Section 5.1. From a purely descriptive perspective, however, we note that the drop in applied-for wages over time appears small compared to the overall variation in our data. While the overall standard deviation in applied-for wages is 13 percent across all monthly observations, Figure 1a shows a month-to-month change of less than 0.1 percent. Comparing regression R^2s reaffirms this conclusion: The regression with individual fixed effects and common duration effects explains 0.811 percent of the variation in average applied-for wages. Dropping the common duration effects, however, only reduces this to 0.810 percent. For describing variation in applied-for wages, level differences across individuals are therefore an order of magnitude more important than changes over time.

Finally, the difference between the red and blue lines in Figure 1 underscores the importance of accounting for dynamic selection. When comparing raw means across completed UI duration, dynamic selection obscures the changes in applied-for wages after month 4 because individuals who apply for permanently higher wages exit unemployment more slowly and gradually make up a larger share of the unemployment pool. We return to this pattern more explicitly in Section 4.4 below.

4.2 Variation over time in other dimensions of search behavior

In Panels B-D of Figure 1, we estimate Equation 6 and 7 for our other main measures of search behavior: the share of applications to full-time jobs, the scope of search and the share of applied-for jobs found via a vacancy posting. We see systematic changes in individual search behavior in all dimensions.

Panels B and C show the results on non-wage job characteristics. After accounting for dynamic selection (red line), the share of applications going to full-time jobs decreases by about 3 percentage points over the first year of unemployment, while the scope of search index increases by about 0.06 standard deviations. We see differences in the timing of these changes however. While the shifts to target more part-time jobs happen gradually over time, changes in scope of search are concentrated entirely during the first three to four months of the unemployment spell. These empirical results lie well beyond the benchmark search model which deals only with wages and predicts no changes in behavior over time. In conjunction with the results on applied-for-wages however, we note that all results point to workers targeting less and less attractive jobs over time, supporting the notion that job seekers are becoming more desperate in their job search (assuming that individuals prefer

Figure 1: Dynamics in applied-for jobs characteristics

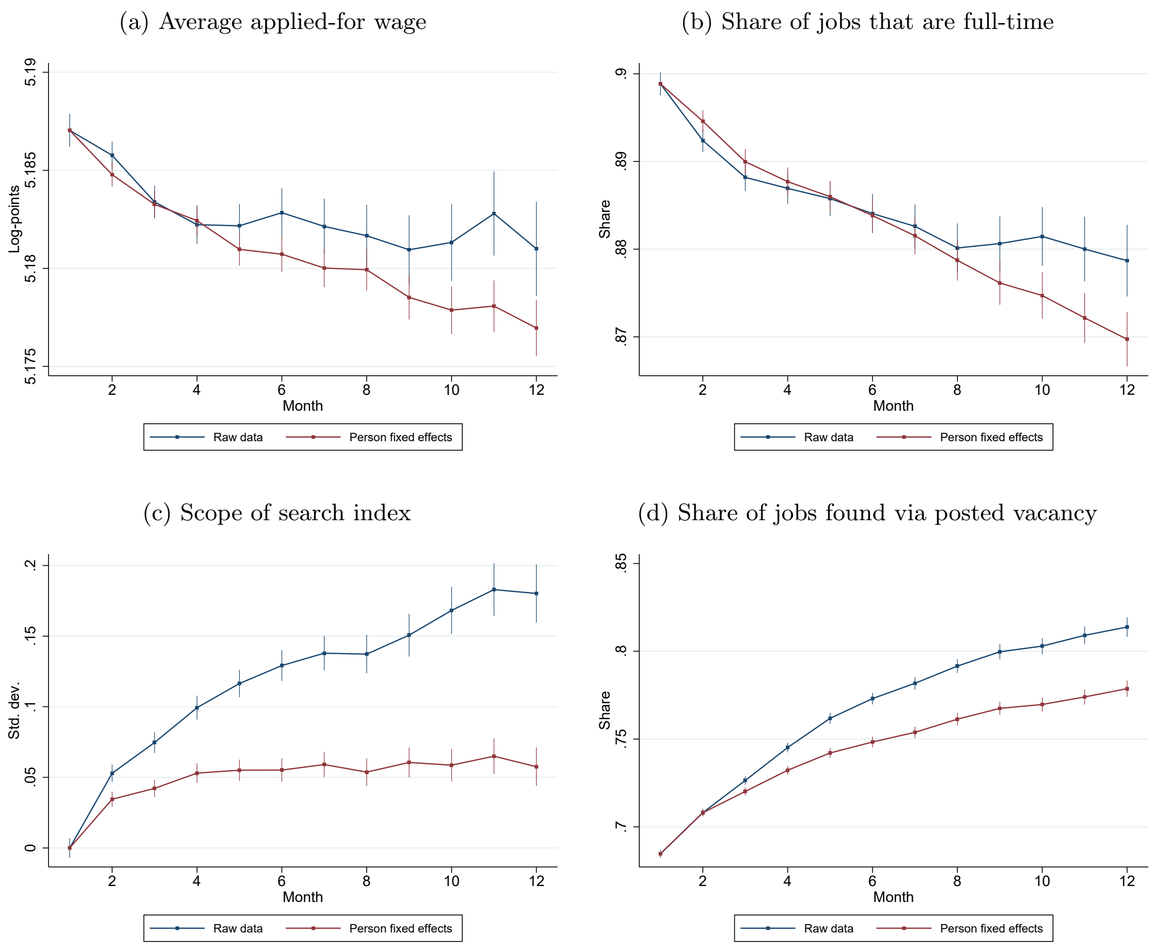

Note: This figure shows the estimates of the duration fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, estimated from Equation 6 (blue line) and Equation 7 (red line). Note that for the red line τ_1 is normalized to the mean of the outcome variable in the first month of unemployment. See Section 3.3 for further details on the measures of applied-for jobs characteristics. Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

full-time jobs, and that industry and occupation switches are costly). We return to this point in Section 5.

In Panel D we examine search channels by looking at the share of applied-for jobs that were found via a vacancy posting. This share increases continuously. After accounting for dynamic selection (red line), the magnitude of this increase is about 10 percentage points over the first year of the unemployment spell. Additional analysis shows that the increased focus on vacancy postings comes both at the expense of applying less to jobs found via social networks and sending less cold contact applications (see Online Appendix C.4 for these results).

In terms of the magnitude of the changes over time, our conclusion on wages carry over to both non-wage job characteristics and search methods: Changes over time are much smaller than level differences across spells (see Table 4).

Finally, across all the panels (B-D), the difference in results with and without individual fixed effects (red vs. blue lines) shows the importance of addressing dynamic selection. Composition changes in the pool of unemployed obscure individual changes in behavior throughout the spell because the pool of unemployed individuals gradually consists of more people who target full-time jobs, have a wider scope of search and rely more on vacancy postings in their job search.

4.3 The fit of the two-way fixed effects model

The two-way fixed effects model used in the preceding sections allows for level differences in search behavior across individuals, but assumes that all individuals adjust behavior in the same way over time. In practice, however, individuals could also exhibit systematic differences in how they adjust their search behavior over time. In addition to being interesting in themselves, such differences in adjustments over time are another source of dynamic selection: some individuals may change their search behavior at a systematically different rate, while also exiting unemployment more slowly leading to dynamic selection *on changes.* Importantly, dynamic selection on changes would not be addressed by the use of individual fixed effects.[11]

In Figure 2, we examine whether such dynamic selection on changes is at play in our data. To do this, we split our analysis sample into three subgroups, depending on the time of UI exit and then re-estimate Equation 7 separately for each group. To reliably compare time paths of behavior, we restrict attention to spells lasting at least 4 months. The panels of the figure plot the estimated duration effects for each of our four measures of search behavior. We continue to normalize the first duration effect to equal the mean in the first month for the relevant subgroup. This implies that gaps across groups in the first month correspond to raw gaps in search behavior across these groups in this month.

If our data exhibit dynamic selection on changes, the adjustment in behavior over time should be

[11] To give an example, assume that individuals who exit unemployment fast exhibit small changes in behavior over time, while individuals who exit unemployment slowly exhibit larger changes throughout the spell. Because the latter group make up a larger and larger share of the sample over time, the estimated duration effects would show search behavior to change slowly initially but faster and faster as the spell progresses. At the individual level, however, changes in behavior are in fact happening at a constant rate over time.

Figure 2: Dynamics in applied-for jobs characteristics by realized duration

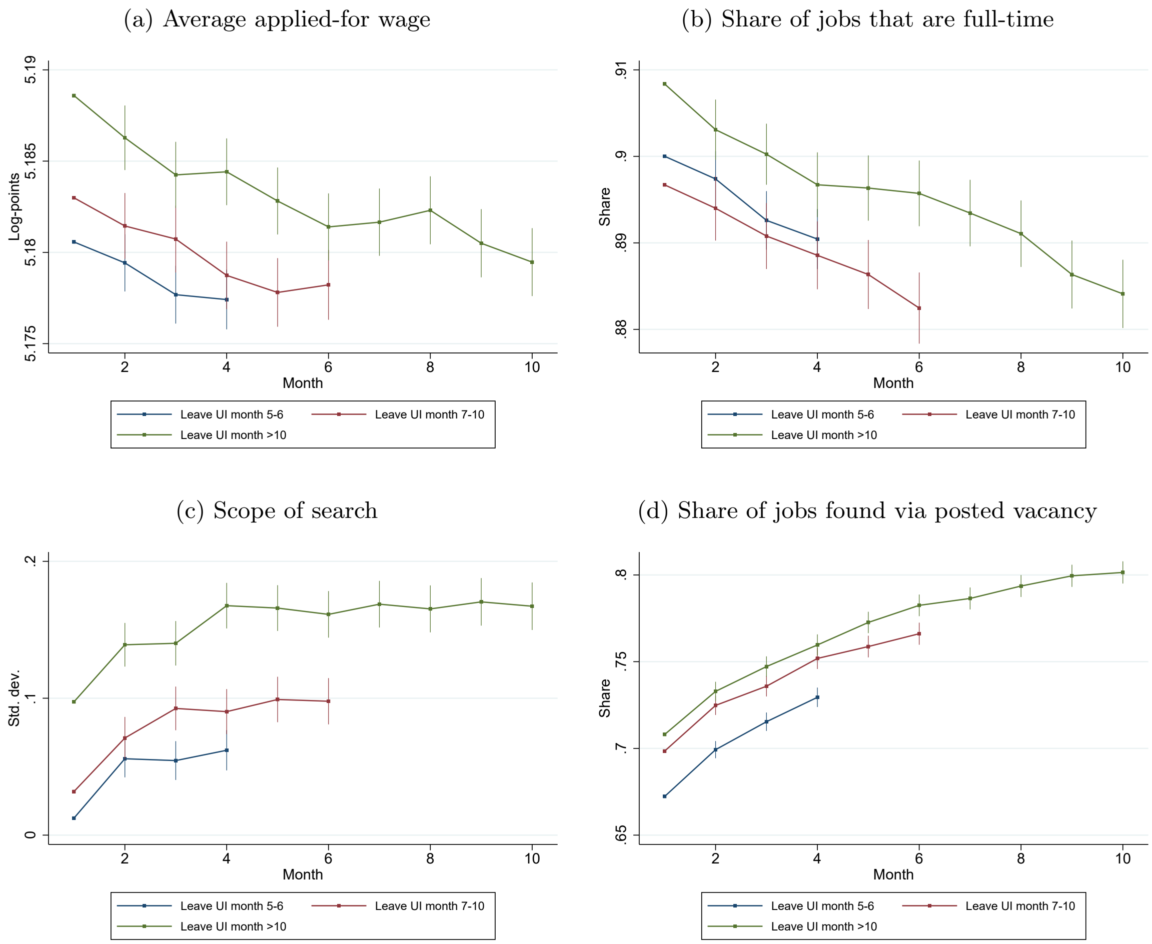

Note: This figure shows the estimates of the duration fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, estimated from Equation 7 (red line) split by groups of individuals with different lengths of realized unemployment duration. Note that τ_1 is normalized to the mean of the outcome variable in the first month of unemployment for the relevant group. See Section 3.3 for further details on the measures of applied-for jobs characteristics. Censored spells with a duration of less than 11 months are excluded from the analysis. Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

different for spells that end sooner. Figure 2 shows no indication of this. While all panels show clear level differences across individuals exiting UI at different times, the time paths of search behavior are remarkably parallel.

More broadly, the results in Figure 2 indicate that the regression model with common duration effects is indeed a good description of the data. To investigate this more systematically, Table 2 compares model fit for the two-way fixed effects model against various competing models, both simpler and more complex. As simpler benchmark models, we consider a regression that only includes a set of observable characteristics for the UI recipient as well as a model that only contains individual fixed effects. The set of observable characteristics includes gender, age, children, education, earnings in the previous job, UI benefit level, previous UI experience, previous industry and previous occupation.[12] As more complex alternatives, we consider two models that allow for differences in how workers change behavior over time: The first model adds interactions between time effects and observable characteristics. The second further adds individual-specific linear time trends. We estimate these different models for our four main measures of job search behavior.

As our measures of fit, we supplement the standard in-sample R^2 with an *out-of-sample* R^2. We compute this out-of-sample R^2 by a four-fold cross-validation procedure that for each individual excludes a fourth of the monthly observations from the estimation and then uses these excluded observations to assess model fit. The out-of-sample R^2 thus measures how well we are able to predict search behavior in a given month of an unemployment spell when given data on search behavior from other months of this spell. Importantly, while in-sample R^2 rises mechanically with model complexity, the out-of-sample R^2 penalizes models with more parameters because of potential overfitting.[13]

Table 2 shows that the data are indeed well described by the two-way fixed effects model with common duration effects. For both applied-for wages and scope of search, the two-way fixed effects model has the single highest out-of-sample fit of all the models considered. For the two remaining outcomes, the model which interacts time effects with observables has a slightly higher out-of-sample R^2; however, the gain relative to the simple two-way fixed effects model is slight. The implication is that - at least to a high degree of approximation - individuals differ markedly in their level of search behavior but adjust behavior in the same way over time. This further implies that initial gaps in search behavior between individuals persist over the entire spell. We return to the theoretical implications of these result in Section 5.4. For empirical work however, the results serve

[12]The exact list of variables are age, log earnings in the previous job, log UI benefits received and dummies for all of the following: being female, having children, education-level (3 levels), education field (1-digit ISCED), previous industry (1-digit) and previous occupation (1-digit). To ensure that the sample does not change across models, we deal with missing values, by setting them to 0 and include missing dummies. Results are robust to using only observations without missing data.

[13]Relative to alternatives like AIC, BIC or Adjusted R^2, the advantage of our out-of-sample R^2 is that it has a natural interpretation as measuring how well we are able to predict search behavior in a given month of an unemployment spell. To ensure that the out-of-sample R^2 is well-defined, and that the sample remains the same across models, we restrict the sample to individuals with at least four months of non-missing data when calculation this statistic. This is also why we opt for four-fold cross validation (more folds would require us to restrict the sample further).

Table 2: Explaining variation in search behavior, model fit

Specification	Applied-for wage R^2_{in}	Applied-for wage R^2_{out}	Full Time R^2_{in}	Full Time R^2_{out}	Scope of search R^2_{in}	Scope of search R^2_{out}	Posted Vacancies R^2_{in}	Posted Vacancies R^2_{out}	Parameters
a) Observables Only	0,357	0,365	0,180	0,183	0,234	0,240	0,108	0,109	45
b) Individual FEs	0,810	0,710	0,695	0,527	0,778	0,651	0,688	0,485	105728
c) Individual and Duration FEs	0,811	**0,711**	0,696	0,528	0,778	**0,652**	0,692	0,491	105739
d) + Observ. X Duration FEs	0,810	0,710	0,697	**0,529**	0,778	0,651	0,693	**0,492**	106220
e) + Individual trend	0,874	0,617	0,808	0,323	0,856	0,560	0,816	0,383	197214

Note: This table reports model fit for different linear regressions models (rows) and different outcome variables (columns). The models are a): Model including individual characteristics: gender, age, children, education, earnings in the previous job, UI benefit level, previous UI experience, previous industry and previous occupation (see footnote 12 for additional details), b): Individual fixed effects only, c): Individual fixed effects and unemployment duration fixed effects, d): Individual fixed effects, unemployment duration fixed effects and interactions between duration dummies and all individual characteristics e): Individual fixed effects, unemployment duration fixed effects, interactions between duration dummies and all individual characteristics, as well as individual-specific linear time trends. The measures of model fit are the standard in-sample R^2 and an out-of-sample R^2obtained from a four-fold cross-validation procedure as described in the text. All models are estimated on our analysis sample of month-by-worker data. When computing the out-of-sample R^2, however, the sample is restricted to unemployment spells that include at least four months of observed search behavior. Imposing this same restriction also for the in-sample R^2 leads to virtually unchanged results. The outcome variables in the regression are the same as in Figure 1: Log average applied-for wage, the share of applied-for jobs that are full time, the scope of search index and the share of applied-for jobs found via a posted vacancy.

as a validation of the parallel trends assumption used in many difference-in-differences analyses of job search.

Finally, two other things are worth noting in Table 2. First, the table reaffirms that changes in behavior over time are orders of magnitudes less important than variation across individuals. Both in and out of sample, the regression models with only individual fixed effects have R^2s in the range of 0.5-0.8 across the different outcomes. Adding the duration fixed effects to these regression increases the R^2s by less than 0.01. Second, in terms of explaining differences in behavior across individuals, the table shows that observable characteristics cannot explain all of the individual variation in applied-for wages; the model with only observable characteristics has an R^2 of 0.36, about half that of the model with individual fixed effects. For our other measures of job search, the role of observables is even more limited.

4.4 Applied-for wages, realized unemployment duration and observables

Results in both Section 4.1 and 4.3 revealed a systematic relationship between applied-for wages and realized unemployment duration. Individuals who systematically target higher-wage jobs exit unemployment at a different rate than others, leading to dynamic selection in Figure 1 and to systematic level differences across unemployment duration groups in Figure 2. We finish our analysis by unpacking this relationship directly and relating it to some standard observables.

In Table 3, we regress UI recipients' average applied-for wage on an indicator for exiting UI very quickly - within 3 months - and an indicator for exiting UI slowly - remaining unemployed for more than 10 months. The omitted category is thus individuals with medium unemployment spells, who exit between months 4 and 10. Given the finding that individuals adjust behavior in parallel over time, we restrict attention to the first month of the unemployment spell here.

Column (1) shows results when only the two indicators are included. It reveals a U-shaped relationship between realized unemployment duration and the average applied-for wage.[14] Relative to those with a medium duration, individuals who exit UI very quickly apply for 0.4 percent higher wages. Individuals with long realized durations apply for even higher wages, however. Relative to the medium group, individuals with long durations apply for 0.7 percent higher wages.

This more complex relationship between applied-for wages and unemployment duration can be rationalized by our benchmark search model with heterogeneous workers from Section 2. If there is a group of workers with a high arrival rate of job offers, the model predicts that these workers will simultaneously exit unemployment fast and target higher-wage jobs. Conversely, if there is a group of workers with a low discount rate or a high value of unemployment, the model predicts that these workers will simultaneously exit unemployment slowly and target higher-wage jobs. The U-shaped relationship seen in Column (1) is thus consistent with having some workers with particularly high job offer arrival rates and other workers with particularly low discount rates or high values of unemployment.

In Columns (2)-(8) we add different UI recipient observables as controls to explore the role of observable characteristics. Column (2) adds basic demographics as controls: gender, age and a dummy for having children. Comparing the coefficient on the duration indicators to those in Column (1), we see that these variables can explain two-thirds of the difference in applied-for wages between the short and medium duration groups. This in particular reflects that males and younger job seekers are more likely to exit unemployment within 3 months and also apply for higher wages (Table 4 in Appendix A.1 shows the relationship between gender, age and unemployment duration).

Columns (3) and (4) instead control for the log of weekly UI benefits or for the log *replacement rate*: weekly UI benefits as a fraction of previous weekly earnings.[15] According to the benchmark search model, individual differences in the level of UI benefits should be a key determinant of search behavior because they directly impact the flow value of unemployment. As predicted by theory, we see a positive relationship between benefits and applied-for wages. Variation in benefits turns out to be quantitatively unimportant for differences in applied-for wages however: in both columns (3) and (4), the R^2 is very small, and the coefficients on the dummies for unemployment duration are unaffected by the added controls. The log replacement rate even enters the regression with a negative coefficient which contradicts the theory.

A likely explanation for the negative coefficient on the replacement rate is that the replacement rate has a mechanical negative relationship with an individual's previous earnings. Previous earnings

[14] Note that this U-shaped relationship is obscured in Figure 2 presented earlier which does not include UI spells shorter than 4 months.

[15] As discussed in Section 3, the benefit level varies across individuals in our setting because UI benefits are determined as 90 percent of previous income up to a certain cap. Since our empirical measure of UI benefits is based on the actual amount of UI benefits paid out , some variation could also reflect e.g. short term holidays.

are likely to correlate positively with applied-for wages if individuals apply to a similar type of jobs as those they previously worked in. In Column (5) we verify that this is indeed the case by controlling for the log of previous earnings. Again, however, the R^2 is very small, and the coefficients on the dummies for unemployment duration are unaffected.

Column (6) shows that education can explain a very substantial part of differences in applied-for wages. Including dummies for having a Upper Secondary or Bachelor/Masters degree increases the R^2 to 0.22 and removes most of the gap in applied-for wages between the medium and long unemployment spell groups. This reflects that more educated individuals apply for markedly higher wages, while simultaneously also having a higher probability of becoming long-term unemployed (Table 4 in Appendix A.1 shows the relationship between educational attainment and unemployment duration). Viewed through the lens of the benchmark model, this would suggest that highly educated individuals have a higher flow value of unemployment or have a lower discount rate. The latter in particular is consistent with standard theories of educational choice.

In Column (7) we examine the role of past experiences with the UI system by including dummies for having previously been on UI within the past one and two years. While past UI experience cannot explain much of the overall variation in search behavior, it can, however, explain a substantial part of the difference between medium and long-term unemployed. This reflects the fact that individuals with past UI experience systematically target lower-paying jobs while simultaneously exiting unemployment faster. This could reflect a learning effect where individuals with recent UI experience are better informed about the prospects of finding a job. We return to learning in Section 5.2.

Finally, in Column (8) we combine the control variables from all the other columns. This reduces the gap in applied-for wages between individuals with short and medium unemployment durations to 0.15 percent and the gap between those with long and medium duration to 0.2 percent higher wages. In sum, most of the relationship between applied-for wages and realized unemployment duration can be explained by observables, but about a third remains unexplained.

For completeness, Online Appendix C.6 repeats the analysis from Table 3 for our other measures of search behavior. For all of these we see monotonic relationships between search behavior and realized unemployment duration. These relationships are largely unrelated to observables.

4.5 Additional results and robustness

The online appendix presents a range of additional results and robustness checks which we briefly discuss here. First, in Online Appendix C.2 we show that the decline in average applied-for wage occurs because the entire distribution shifts downward, rather than because of changes occurring only at the bottom of the distribution. As we expand on in the appendix, this provides some evidence that workers target their search towards specific wage-levels, rather than searching randomly and rejecting job opportunities below a certain cutoff.

Second, in Online Appendix C.3 Figure 8, we report results separately for the three components of the scope of search index (relatedness of applied-for occupation to the previous job, relatedness

Table 3: Applied-for wages, realized unemployment duration and observables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Duration: 2-3 months	0.00369***	0.00134	0.00270***	0.00384***	0.00376***	0.00409***	0.00419***	0.00145*
	(0.00102)	*(0.00100)*	*(0.00103)*	*(0.00102)*	*(0.00102)*	*(0.000901)*	*(0.00102)*	*(0.000866)*
Duration: >10 months	0.00723***	0.00885***	0.00748***	0.00725***	0.00728***	0.00260**	0.00514***	0.00215*
	(0.00148)	(0.00145)	*(0.00148)*	*(0.00148)*	*(0.00148)*	*(0.00131)*	*(0.00148)*	(0.00125)
Female		-0.0539***						-0.0670***
		(0.00104)						*(0.00091)*
Age		-0.000734***						-0.000113***
		(0.00004)						*(0.00004)*
One or more children		-0.00101						-0.00466***
		(0.00127)						(0.00110)
Log UI payout			0.0246***					0.0292***
			(0.00286)					*(0.00258)*
Log UI payout/previous earnings				-0.00236***				-
				(0.00050)				-
Log previous earnings					0.00286***			0.0107***
					(0.000479)			*(0.00047)*
Upper secondary						0.0182***		0.0224***
						(0.00129)		*(0.00125)*
Short tertiary/Bachelor/Master/Doctoral						0.142***		0.154***
						(0.00130)		(0.00128)
Any UI experience past year							-0.0338***	-0.0148***
							(0.00212)	*(0.00181)*
Any UI experience past two years							-0.0171***	-0.0116***
							(0.00143)	*(0.00121)*
Constant	5.182***	5.240***	4.979***	5.181***	5.157***	5.115***	5.188***	4.821***
	(0.00073)	*(0.00172)*	*(0.02360)*	*(0.00075)*	*(0.00417)*	*(0.00123)*	*(0.00076)*	*(0.02040)*
Observations	81,302	81,302	81,302	81,302	81,302	81,302	81,302	81,302
R-squared	0.000	0.044	0.001	0.001	0.001	0.221	0.011	0.296

Note: This table reports estimates from regressions where the dependent variable is the average applied for wages in the first month of unemployment. In all regressions we include dummy variables distinguishing spells by their “completed” unemployment duration. We distinguish between short, medium and long unemployment spells (medium spells are the omitted category). Short spells are 3 months or less, long spells are more than 10 months. Censored spells with a duration of less than 11 months are excluded from the analysis. Across the different columns, we change the set of control variables to assess the importance of these variables in explaining the differences across unemployment duration groups. For further information on control variables, see Online Appendix B.2. Note that the log replacement rate drops out in column 8 because of perfect collinearity with log UI payout and log previous earnings. *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

of applied-for industry, and commuting time), as well as for alternative definitions of industry and occupation-relatedness. Results all mirrors those found for the overall scope of search index in Figure 1c above.

Finally, in Online Appendix C.4 we present results for all the different search and application channels on which we have information. As already discussed, the increased focus on vacancy postings in Figure 1d comes at the expense of applying less to jobs found via social networks and fewer cold contact applications. In Online Appendix C.5 we further show that the changes in applied-for jobs shown in Figures 1a-1c do not appear to stem from a composition effect due to this change in search methods.

5 Implications for Theory

The empirical analyses in the previous section revealed several results that cannot be rationalized within the benchmark search model from Section 2. In this section we discuss the implications of our results for richer theories of job search.

5.1 The quantitative importance of declining applied-for wages

In contrast to the predictions of the benchmark model, our results show that applied-for wages decline 1 percent over the first year of an unemployment spell. As discussed in Section 4.1, this drop in applied-for wages over time is negligible relative to the variation in applied-for wages across individuals. A question remains, however, about whether this decline is economically important.

A natural focal point is to ask how much the observed 1 percent drop in applied-for wages affects the hazard rate out of unemployment. In the notation from Section 2, this is measured by the elasticity of the hazard rate, h, with respect to the average-applied-for wage, w^*.[16] It turns out that the benchmark random job search framework imply a formula for this elasticity under weak conditions. As long as the reservation wage is binding (i.e. w^R is inside the support of F) and the wage distribution is continuous, it follows from the definitions of the hazard rate (Equation 3) and the average applied-for wage (Equation 5) that the elasticity is a simple function of the average applied-for wage and the reservation wage, w^R

$$\eta^h_{w^*} = -\frac{1}{1 - \frac{w^R}{w^*}} \tag{8}$$

In Appendix A.3 we provide a full derivation and additional discussion of this formula. We make two brief remarks here. First, because the formula hinges only on the definition of the hazard rate and the average applied-for wage, it applies very generally. For example, it requires no assumptions

[16]While our benchmark model cannot explain the within spell changes in applied-for wages we use the benchmark model as a reference framework to quantify the economic effects of the changes we observe. In other words we ask: if we were to decrease reservation wages permanently by 1 percent at the beginning of the unemployment spell, what would the effect on job finding be? Note that for a given wage change we should expect the change in the value of unemployment to be larger the smaller the elasticity(in absolute terms). A small elasticity suggests that workers were willing to take wage cuts although chances of finding employment only increased marginally.

about how job seekers actually choose their reservation wage (and thus their average applied-for wage). Second, the key intuition is that when w^R and w^* are far apart, a given change in the average applied-for wage implies a smaller change in the reservation wage and thus also implies a smaller change in the set of jobs that the job seekers apply to (ceteris paribus). Accordingly, when w^R and $w^*$are far apart, the formula implies that the job finding elasticity will be small (and vice versa if they are close).

Importantly, the formula in Equation 8 allows us to gauge the job finding elasticity from our data on applied-for wages. In particular, we use our ability to observe several applied-for wages within individuals over time to construct the ratio of the lowest applied-for wage to the average applied-for wage *for each spell*. This exercise exploits the longitudinal dimension of our data and gives us an estimate of $\frac{w^R}{w^*}$, which we can use to compute $\eta^h_{w^*}$ for each spell. Depending on implementation details, we arrive at a median elasticity between 4.1 and 7.7 (see Appendix A.3 for details). Relative to a counterfactual where applied-for wages were constant, the fact that applied-for wages drop by 1 percent thus means that the hazard rate is 4.1-7.7 percent higher one year into the unemployment spell. For research that aims to understand hazard rates, this is a substantial effect. As a benchmark, it corresponds to the total effect of reducing UI benefits by 8-15 percent based on recent estimates.[17]

5.2 Introducing declines in applied-for wages into the benchmark model

The previous section suggests that it is quantitatively important to match the decline in applied-for wages if one wants to understand the evolution of the job finding hazard over time. Our results thus point to the necessity of including non-stationarity in job search models.

A simple way to introduce non-stationarity in our benchmark model is to allow for anticipated time variation in parameters. This will generate changes in search behavior both directly at the time when parameters change, but also in expectation of future changes. An empirically relevant example is to assume that the flow value of unemployment drops when UI benefits expire. Because workers are forward-looking, this implies that the continuation value of unemployment gradually decreases until benefits expire, leading to a continuous drop in applied-for wages (Mortensen (1977)). Other common extensions include a declining job offer arrival rate due to stigma, human capital depreciation or some other kind of screening on the employer side (see e.g. Lockwood, 1991; Pissarides, 1992; Acemoglu, 1995; Kroft et al., 2013) or a gradually lower flow value of unemployment as workers deplete their savings and have lower consumption (see e.g. Lentz and Tranæs, 2005; Mortensen, 1986). Any of these extensions introduce declines in applied-for wages over time as seen in the data.

Another way to introduce time varying behavior is via learning and changing beliefs (see e.g. Spinnewijn, 2015; Conlon et al., 2018; Mueller et al., 2021). When workers have different beliefs about job prospects and the search process, this is a source of variation in search behavior across

[17] Schmieder and von Wachter (2016) survey the estimated effects of changes in UI benefit levels on the unemployment durations. Their median estimate corresponds to an elasticity of the hazard rate of -0.53 respectively (using a constant hazard approximation). Based on this, benefits would thus need to be lowered by 7.7 to 14.5 percent to induce increase in the hazard of between 4.1 and 7.7 percent.

workers (i.e. differences in levels). Moreover, if workers update their beliefs and learn from their experiences while unemployed, this will generate variation in behavior over time. This is because each additional week of unsuccessful job search is likely to be a signal that employment prospects are worse than expected (see e.g. Burdett and Vishwanath, 1988; Gonzalez and Shi, 2010). The fact that workers with previous UI experience target systematically lower-paying jobs and exit unemployment faster in our data can be seen as supporting this idea (See Table 3 and Table 4).

Behavioral biases may also play a role in generating the time patterns of search behavior that we see in our data. Recent work has emphasized models of loss-aversion and reference-dependence for explaining time patterns of search behavior (see e.g. Dellavigna et al., 2017; DellaVigna et al., 2021). Brown et al. (2011) present an overview of other behavioral biases which could also explain declining reservation wages throughout the unemployment spell such as e.g., the sunk cost fallacy or subjective search costs. It is also possible that job seekers simply become more discouraged or feel more stigmatized over time in unemployment (see e.g. Krueger et al. (2014); Krueger and Mueller (2011); Frijters and van der Klaauw (2006)).

As we return to further below, time variation in behavior may also be explained by so-called stock-flow matching.

5.3 Generating variation in other dimensions of search behavior

Besides applied-for wages, our results show substantial variation in other dimensions of search behavior both across individuals and over time. Matching these requires more substantial deviations from the benchmark model which only deals with wages.

We find substantial variation in applied-for hours and the scope of search, as measured by the implied commute and the relatedness of the job's occupation and industries to that of the worker's previous job. One way to generate this is to extend the benchmark model so that jobs differ also in non-wage characteristics such as hours, commuting time and whether they require the worker to learn a new occupation/industry. When these characteristics impact the utility of the job, job search decisions are no longer described by a reservation wage, but instead by a total reservation utility that the job must provide given all its characteristics (see e.g. Hall and Mueller, 2018; Le Barbanchon et al., 2020). Since workers only apply for jobs above their own reservation utility, this will generate variation in both hours and the scope of search. Moreover, if the model is non-stationary with a decreasing value of unemployment, the reservation utility will also be decreasing over time. This is consistent with the pattern that workers seem to target worse jobs over time in all dimensions: lower applied-for wages, fewer applications to full-time jobs and a broader scope of search.

In our data, we also observe substantial variation in search channels both across individuals and over time. Over time, workers rely more and more on formally posted vacancies as opposed to informal search methods such as social networks. Given the evidence that informal channels are more effective at job finding (Hensvik and Skans (2016); Dustmann et al. (2016)), a possible interpretation is that workers direct attention to the most attractive search channels first, but gradually exhaust these over the course of an unemployment spell.

A seminal model capturing this idea is the systematic (sequential) search model of Salop (1973). The empirical pattern can also be interpreted within the stock-flow matching framework (see e.g. Coles and Smith, 1998; Ebrahimy and Shimer, 2010; Coles and Petrongolo, 2008). Here newly unemployed workers face a large potential stock of vacancies or job opportunities. Over time, however, these opportunities are exhausted, and workers who remain unemployed are limited to waiting for new job opportunities to arise. In Online Appendix C.3.2 we present some empirical evidence in favor of the stock-flow mechanism: using a rough proxy for newly posted vacancies, we indeed find that UI seekers rely more and more on newly posted vacancies over the course of an unemployment spell as the theory predicts. The implication is that workers who have been unemployed for longer may have different search behavior because they have exhausted more potential job opportunities and search channels. We note that this type of model also provides an alternative interpretation for the time variation in scope of search: if workers first exhaust job options nearby and in related industries/occupations, we should see the scope of search gradually broaden over the unemployment spell.[18]

5.4 Implications of the two-way fixed effects model

A final main takeaway from our empirical analysis is that variation in search behavior is very well described by a two-way fixed effect model with individual fixed effects and common duration effects (Equation 7). Moreover, within this model the individual fixed effects explain most of the variation.

In addition to highlighting the importance of having heterogeneity both across workers and over time, this result also places useful restrictions on the nature of this heterogeneity. In particular, while worker heterogeneity should create large level differences in (log) applied-for wages, it should not lead to differential time paths. As an example of the restrictions this impose considera learning model where worker heterogeneity purely reflects differences in beliefs. In this model, our results rules out the possibility of having both optimistic and pessimistic workers who converge towards correct beliefs over time because such convergence would imply differential changes in search behavior for the two groups.[19]

Conversely, an example of a common modeling approach that is consistent with our results is the use of proportional value functions: In some search models, heterogeneity across workers is introduced via an ability parameter that directly scales the equilibrium value of unemployment (for recent examples see e.g. Hall and Mueller, 2018; Flinn and Mullins, 2021). Combined with a specific assumption on the wage distribution, this approach exactly implies a two-way fixed effect model for the average applied-for wage.

[18] More broadly this relates to models like Moscarini (2001) where workers choose to search in different markets based on their comparative advantage and the job prospects in each market (see also Decreuse (2008); Papageorgiou (2014)).

[19] Of course, in a richer model with multiple sources of heterogeneity the effect of converging beliefs may be offset by other dynamics so that behavior on net still fits the two-way fixed effect model. Our results thus do not rule out the co-existence of optimistic and pessimistic workers with converging beliefs in richer models.

To see this, we adapt the notation of our benchmark model slightly and let U_{it} be the value of unemployment for worker i at time t in her unemployment spell. Now we assume that worker heterogeneity can be summarized by each worker having a scalar ability level $a_i \geq 1$. Letting $\tilde{U}_t$ denote the value of unemployment for a worker with $a_i = 1$ at time t during their unemployment spell, it is then possible to set up model primitives such that any worker's value of unemployment is proportional to her ability level (Online Appendix D.1 provides an example):

$$U_{it} = a_i \tilde{U}_t \tag{9}$$

The typical advantage of this setup is that it simplifies the model solution, however, as it turns out it can also rationalize the two-way fixed effect model of applied-for wages. To see this, note that if reservation wages and the average applied-for wage are as in the benchmark model (Equations 1 and 5), and if the wage offer distribution is assumed to be Pareto with some shape parameter $\gamma > 1$, the average applied for wage will satisfy:

$$w_{it}^* = \frac{\gamma}{\gamma - 1} r U_{it}$$

Combining this with Equation 9 and taking logs then exactly yields a two-way fixed effects model for the average applied-for wage:

$$log\left(w_{it}^*\right) = \alpha_i + \tau_t \qquad \alpha_i \equiv log\left(a_i\right), \tau_t \equiv log\left(\frac{\gamma}{\gamma - 1} r \tilde{U}_t\right)$$

In this sense, our results provide empirical support for modeling worker heterogeneity as in Equation 9.

6 Conclusion

In this paper, we use new linked administrative data on applied-for jobs to study variation in job search behavior across individuals and over time. We establish two main sets of results.

First, we find large level differences in applied-for wages across individuals. Over time, however, all job seekers adjust their wages downward in the same way during unemployment. Over the first year of unemployment, the average applied-for wage declines by about 1 percent. Relative to the overall variation in applied-for wages, this is a small change, however, we show that it is quantitatively important for the job finding rate in standard models. This underscores the importance of including both extensive worker heterogeneity and non-stationarity in job search models. The fact that workers adjust behavior in the same way over time implies theoretical restrictions on how dynamics and worker heterogeneity interact. Empirically, it also serves as a validation of the parallel trends assumption used in many difference-in-difference analyses of job search.

Second, in addition to wages, we also find extensive variation both in the non-wage characteristics of the jobs that workers apply for and in the search channels that workers use. This variation follows patterns similar to those for wages. In particular, all workers adjust behavior in a similar way over time across these dimensions, and the adjustments are consistent with workers targeting less and less attractive jobs. These results suggest that search decisions about non-wage job characteristics and search channels warrant more attention in both theoretical and empirical work on job search.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333. 3.3, B.4

Acemoglu, D. (1995). Public Policy in a Model of Long-term Unemployment. *Economica*, 62(246):161. 5.2

Ahrens, A., Hansen, C. B., and Schaffer, M. E. (2019a). Lassopack: Model Selection and Prediction with Regularized Regression in Stata. *IZA Discussion Papers*. B.4

Ahrens, A., Hansen, C. B., and Schaffer, M. E. (2019b). LASSOPACK: Stata module for lasso, square-root lasso, elastic net, ridge, adaptive lasso estimation and cross-validation. *Statistical Software Components*. B.4

Allen, M. T., Waugh, G., Shaw, M., Tsacoumis, S., Rivkin, D., Lewis, P., Brendle, M., Craven, D., Gregory, C., and Connell, D. (2012). The Development and Evaluation of a New ONET Related Occupations Matrix Volume I : Report. *National Center for ONET Development*, I. B.5

Altmann, S., Cairo, S., Mahlstedt, R., and Sebald, A. (2019). Complexity and the Effectiveness of Public Policy. *Manuscript*. B.6.1

Altmann, S., Falk, A., Jäger, S., and Zimmermann, F. (2018). Learning about job search: A field experiment with job seekers in Germany. *Journal of Public Economics*, 164:33–49. 1

Banfi, S., Choi, S., and Villena-Roldan, B. (2019a). Deconstructing Job Search Behavior. *Working Paper*. 1

Banfi, S., Choi, S., and Villena-Roldán, B. (2019b). Sorting On-line and On-time. *University of Bristol Discussion Paper 19/706*. 1

Banfi, S. and Villena-Roldán, B. (2019). Do high-wage jobs attract more applicants? Directed search evidence from the online labor market. *Journal of Labor Economics*, 37(3):715–746. 1

Belloni, A., Chernozhukov, V., and Hansen, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *Review of Economic Studies*, 81(2):608–650. B.4

Belloni, A., Chernozhukov, V., Hansen, C., and Kozbur, D. (2016). Inference in High-Dimensional Panel Models With an Application to Gun Control. *Journal of Business and Economic Statistics*, 34(4):590–605. B.4

Belot, M., Kircher, P., and Muller, P. (2019). Providing advice to jobseekers at low cost: An experimental study on online advice. *Review of Economic Studies*, 86(4):1411–1447. 1

Brown, M., Flinn, C. J., and Schotter, A. (2011). Real-time search in the laboratory and the market. *American Economic Review*, 101(2):948–974. 5.2

Burdett, K. and Vishwanath, T. (1988). Declining reservation wages and learning. *Review of Economic Studies*, 55(4):655–666. 5.2

Card, D., Chetty, R. A., and Weber, A. (2007). The spike at benefit exhaustion: Leaving the unemployment system or starting a new job? *American Economic Review*, 97(2):113–118. 3.2

Christensen, B. J. and Kiefer, N. M. (1991). The exact likelihood function for an empirical job search model. *Econometric Theory*, 7(4):464–486. 23

Coles, M. and Petrongolo, B. (2008). A test between stock-flow matching and the random matching function approach. *International Economic Review*, 49(4):1113–1141. 5.3

Coles, M. G. and Smith, E. (1998). Marketplaces and Matching. *International Economic Review*, 39(1):239. 5.3

Conlon, J., Pilossoph, L., Wiswall, M., and Zafar, B. (2018). Labor Market Search With Imperfect Information and Learning. *National Bureau of Economic Research Working Paper Series*. 5.2

Cox, J. and Oaxaca, R. (1990). Unemployment Insurance and Job Search. *Research in Labor Economics*, 11. 1

Danish Economic Council (2014). Danish Economy, Autumn Report. Technical report. 4, 3.2

Decreuse, B. (2008). CHOOSY SEARCH AND THE MISMATCH OF TALENTS*. *International Economic Review*, 49(3):1067–1089. 18

DellaVigna, S., Heining, J., Schmieder, J. F., and Trenkle, S. (2021). Evidence on Job Search Models from a Survey of Unemployed Workers in Germany. *SSRN Electronic Journal*. 1, 3.2, 3.4, 5.2

Dellavigna, S., Lindner, A., Reizer, B., and Schmieder, J. F. (2017). Reference-dependent job search: Evidence from Hungary. *Quarterly Journal of Economics*, 132(4):1969–2018. 5.2

Dustmann, C., Glitz, A., Schonberg, U., and Brucker, H. (2016). Referral-based Job Search Networks. *The Review of Economic Studies*, 83(2):514–546. 5.3

Ebrahimy, E. and Shimer, R. (2010). Stock-flow matching. *Journal of Economic Theory*, 145(4):1325–1353. 5.3

Engmann, T. S. and Weiskopf, S. V. (2019). Mænd får oftere job gennem netværk og bliver lidt oftere headhuntet end kvinder. *Statistics Denmark*. B.6.2

Faberman, J., Mueller, A. I., Sahin, A., and Topa, G. (2021). Job Search Behavior Among the Employed and Non-Employed. *SSRN Electronic Journal*. 3.4

Faberman, R. J. and Kudlyak, M. (2019). The intensity of job search and search duration. *American Economic Journal: Macroeconomics*, 11(3):327–357. 1, 3.4, B.2

Flinn, C. and Mullins, J. (2021). Firms' Choices of Wage-Setting Protocols. *Working Paper*. 5.4

Fluchtmann, J., Glenny, A. M., Harmon, N., and Maibom, J. (2019). Gender gaps in job search and hiring outcomes. *In: Essays on the Economics of Job Search (PhD dissertation)*. 1, 3, B.4, B.6

Frijters, P. and van der Klaauw, B. (2006). Job Search with Nonparticipation. *The Economic Journal*, 116(508):45–83. 5.2

Ganong, P. and Noel, P. (2019). Consumer spending during unemployment: Positive and normative implications. *American Economic Review*, 109(7):2383–2424. 3.2

Garibaldi, P., Moen, E. R., and Sommervoll, D. E. (2016). Competitive on-the-job search. *Review of Economic Dynamics*, 19:88–107. 52

Gee, L. K. (2019). The more you know: Information effects on job application rates in a large field experiment. *Management Science*, 65(5):2077–2094. 1

Gonzalez, F. M. and Shi, S. (2010). An Equilibrium Theory of Learning, Search, and Wages. *Econometrica*, 78(2):509–537. 5.2

Hall, R. E. and Mueller, A. I. (2018). Wage dispersion and search behavior: The importance of nonwage job values. *Journal of Political Economy*, 126(4). 1, 5.3, 5.4

Harmon, N. a. (2013). Are Workers Better Matched in Large Labor Markets ? *Working Paper*. 3.3, B.5

Hensvik, L. and Skans, O. N. (2016). Social networks, employee selection, and labor market outcomes. *Journal of Labor Economics*, 34(4):825–867. 5.3

Hornstein, A., Krusell, P., and Violante, G. L. (2011). Frictional Wage Dispersion in Search Models: A Quantitative Assessment. *American Economic Review*, 101(7):2873–98. 24

Kesternich, I., Schumacher, H., Siflinger, B., and Valder, F. (2022). Reservation wages and labor supply. *Journal of Economic Behavior & Organization*, 194:583–607. 1

Kroft, K., Lange, F., and Notowidigdo, M. J. (2013). Duration dependence and labor market conditions: Evidence from a field experiment. *Quarterly Journal of Economics*, 128(3). 5.2

Krueger, A. B., Cramer, J., and Cho, D. (2014). Are the long-term unemployed on the margins of the labor market? *Brookings Papers on Economic Activity*, Spring 201:229–299. 5.2

Krueger, A. B. and Mueller, A. (2011). Job search, emotional well-being, and job finding in a period of mass unemployment: Evidence from high-frequency longitudinal data. *Brookings Papers on Economic Activity*, (1):1–57. 1, 5.2

Krueger, A. B. and Mueller, A. I. (2016). A contribution to the empirics of reservation wages. *American Economic Journal: Economic Policy*, 8(1):142–179. 1, 1

Kudlyak, M., Lkhagvasuren, D., and Sysuyev, R. (2013). Systematic Job Search: New Evidence from Individual Job Application Data. *SSRN Electronic Journal*. 1

Kuhn, P. and Shen, K. (2013). Gender discrimination in job ads: Evidence from china. *Quarterly Journal of Economics*, 128(1):287–336. 1

Le Barbanchon, T., Rathelot, R., and Roulet, A. (2019). Unemployment insurance and reservation wages: Evidence from administrative data. *Journal of Public Economics*, 171:1–17. 1

Le Barbanchon, T., Rathelot, R., and Roulet, A. (2020). Gender Differences in Job Search: Trading off commute against wage. *The Quarterly Journal of Economics*, 136(1):381–426. 1, 5.3

Lentz, R., Maibom, J., and Moen, E. (2021). Competitive or Random Search? 52

Lentz, R. and Tranæs, T. (2005). Job search and savings: Wealth effects and duration dependence. *Journal of Labor Economics*, 23(3):467–489. 5.2

Lichter, A. and Schiprowski, A. (2021). Benefit duration, job search behavior and re-employment. *Journal of Public Economics*, 193:104326. 1

Lockwood, B. (1991). Information Externalities in the Labour Market and the Duration of Unemployment. *The Review of Economic Studies*, 58(4):733. 5.2

Marinescu, I. (2017). Job search monitoring and assistance for the unemployed. *IZA World of Labor*. 3.4, 8, B.1

Marinescu, I. and Rathelot, R. (2018). Mismatch unemployment and the geography of job search. *American Economic Journal: Macroeconomics*, 10(3):42–70. 1

Marinescu, I. and Skandalis, D. (2021). Unemployment insurance and job search behavior. *Quarterly Journal of Economics*, 136(2):887–931. 1, 2.1, 3.2, 3.4, A.2

McCall, J. J. (1970). Economics of Information and Job Search. *The Quarterly Journal of Economics*, 84(1):113. 2

McVicar, D. (2020). The impact of monitoring and sanctioning on unemployment exit and job-finding rates. *IZA World of Labor*. 3.4, B.1

Moen, E. R. (1997). Competitive search equilibrium. *Journal of Political Economy*, 105(2):385–411. 20

Mortensen, D. T. (1977). Unemployment Insurance and Job Search Decisions. *Industrial and Labor Relations Review*, 30(4):505. 2, 5.2

Mortensen, D. T. (1986). Job search and labor market analysis. *Handbook of Labor Economics*, 2(C):849–919. 5.2

Mortensen, D. T. and Pissarides, C. A. (1999). Chapter 39 New developments in models of search in the labor market. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 3 PART of *Handbook of Labor Economics*, chapter 39, pages 2567–2627. Elsevier. A.2

Moscarini, G. (2001). Excess Worker Reallocation. *The Review of Economic Studies*, 68(3):593–612. 18

Mueller, A. I., Spinnewijn, J., and Topa, G. (2021). Job Seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias. *American Economic Review*, 111(1):324–363. 1, 5.2

Neffke, F. and Henning, M. (2013). Skill relatedness and firm diversification. *Strategic Management Journal*, 34(3):297–316. B.5

Nekoei, A. and Weber, A. (2017). Does extending unemployment benefits improve job quality? *American Economic Review*, 107(2):527–561. 2.1, A.2, C.2

Papageorgiou, T. (2014). Learning Your Comparative Advantages. *The Review of Economic Studies*, 81(3):1263–1295. 18

Pissarides, C. A. (1992). Loss of Skill During Unemployment and the Persistence of Employment Shocks. *The Quarterly Journal of Economics*, 107(4):1371–1391. 5.2

Rogerson, R., Shimer, R., and Wright, R. (2005). Search-theoretic models of the labor market: A survey. *Journal of Economic Literature*, 43(4):959–988. 2, 2, 2.2, A.2

Salop, S. C. (1973). Systematic Job Search and Unemployment. *The Review of Economic Studies*, 40(2):191. 5.3

Schmieder, J. F. and von Wachter, T. (2016). The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation. *Annual Review of Economics*. 17

Shimer, R. (2005). The assignment of workers to jobs in an economy with coordination frictions. *Journal of Political Economy*, 113(5):996–1025. 20

Skandalis, D. (2018). Breaking News: The Role of Information in Job Search and Matching. *Mimeo*. 1

Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association*, 13(1):130–167. 5.2

Wright, R., Kircher, P., Julien, B., and Guerrieri, V. (2020). Directed search: A guided tour. *Journal of Economic Literature*, 23884:1–120. 20

Part I

Appendix (in paper)

A Appendix

A.1 Sample statistics

Table 4: Summary statistics - Final sample

		By Realized Unemployment Duration				
	Total	< 4 Months	4-6 Months	7-10 Months	> 10 Months	Censored†
Demographics						
Female	0.53	0.51	0.53	0.55	0.54	0.55
Age	38.12	37.10	38.27	38.50	40.08	38.38
One or more children	0.203	0.20	0.21	0.21	0.19	0.21
Education††						
Primary/lower secondary	0.13	0.13	0.15	0.141	0.144	0.12
Upper secondary	0.43	0.461	0.442	0.42	0.39	0.36
Short tertiary/Bachelor/Master/Doctoral	0.44	0.41	0.40	0.43	0.46	0.52
Labor history						
Weekly UI payout	4,086	4,007	3,814	3,796	3,767	5,033
Previous earnings (weekly)	7,439	7,058	7,351	7,453	8,010	7,935
Any UI experience past year	0.07	0.09	0.09	0.05	0.04	0.02
Any UI experience past two years	0.18	0.20	0.20	0.17	0.14	0.13
Spell						
Average spell length (weeks)	24.81	11.47	22.00	36.36	64.66	20.45
Applied-for jobs per week	1.58	1.52	1.60	1.61	1.66	1.61
Observed spell ends with employment	0.66	0.89	0.86	0.83	0.43 (0.43)†††	.
N (number of spells)	127,695	45,831	29,801	14,320	15,708	22,035

Notes: † Censored spells are defined as spells which are censored with duration less than 11 months ††Education: Degree level of highest completed education. ††† The number in parenthesis gives the share of spells which are above 10 months but right-censored in our panel.

A.2 Linking application data to the benchmark search model

In the standard interpretation of the benchmark search model, job applications are entirely absent: some exogenous process simply leads firms to contact workers and extend job offers at some exogenous rate. As is well-known however, job offers in standard models can equivalently be interpreted as stemming from a process where workers contact firms instead (see for example Mortensen and Pissarides, 1999). As we show below, an interpretation in this vein provides a straightforward way to relate the model to data on applied-for jobs.

Instead of job offers arriving at some rate, we assume that at some exogenous rate, $\frac{\lambda}{p}$, unemployed workers learn about a potential job, drawn at random from the job offer distribution. Here $\lambda \geq 0$ and $1 \geq p \geq 0$. Upon learning about the job, they observe the associated wage $w \sim F(w)$ and then decide whether to apply for the job or not. If they apply for the job, there is a fixed probability p that the application is successful and they get the job, otherwise they stay unemployed.

Now let $A(w)$ be the continuation value of sending an application to a job paying w. Standard arguments give the following continuous time Bellman equation for an optimizing worker:

$$rU = b + \frac{\lambda}{p} \int_0^\infty \max \{A(w) - U, 0\} \, dF(w) \tag{10}$$

Now when a worker sends an application to a job paying w, there is a probability $(1 - p)$ of staying unemployed and a probability p of getting the job which results in a continuation value of $\frac{w}{r}$. Accordingly $A(w)$ satisfies

$$A(w) = p \cdot \frac{w}{r} + (1 - p) \cdot U$$

Plugging this into equation 10 yields the standard Bellman equation from the benchmark model (Rogerson et al. (2005)):

$$rU = b + \lambda \int_0^\infty \max \left\{ \frac{w}{r} - U, 0 \right\} dF(w) \tag{11}$$

Optimal worker behavior and the hazard rate are thus still described by Equations 1, 2 and 3 from the main text. The only difference is in the interpretation that the reservation wage, w^R, which now represents both the lowest wage the worker would accept if offered but also the lowest wage that the worker applies for.

Although the model and its predictions are unchanged, the advantage of the interpretation above is that it allows us to think about data on job applications using Equation 4: since potential jobs that workers encounter are random draws from the job offer distribution F, and since workers only apply to jobs offering a wage $w \geq w^R$, it follows that the wages of applied-for-jobs are random draws from the part of F that lies above w^R. This is our preferred way to relate the application data to theory.

For completeness, we note that another possible way to link application data to the theory is to start from the partial equilibrium directed search model of Nekoei and Weber (2017). Instead of workers choosing a reservation wage in the face of a fixed offer arrival rate and offer distribution, this model posits that workers choose a *target wage,* but face the trade-off that targeting a higher wage leads to a lower rate of job finding.[20] It is then natural to equate the average applied-for wage in application data with the chosen target wage. Under weak conditions, the partial equilibrium directed search model turns out to deliver similar predictions as our benchmark model with respect to the average applied-for wage, w^*, satisfying the same equation that we rely on (Equation 5).

[20]In the literature of competitive (directed) job search this relationship arises in equilibrium, see e.g. Moen (1997); Wright et al. (2020); Shimer (2005).

Additional details and discussions are provided in Nekoei and Weber (2017) and Marinescu and Skandalis (2021).

A.3 The effect of changes in applied-for-wages on job finding

Our empirical results show that workers systematically lower their average applied-for wage over time. An obvious question is how much such changes impact job finding. Standard random search models such as our benchmark model from Section 2 have strong implications here.

Specifically, consider a worker facing an arrival rate λ, a wage offer distribution, F, and whose search behavior involves an (initial) reservation wage of w^R. When combined with job application behavior as in Section 2, this translates into an average applied-for wage of w^* as defined by Equation 5 and a hazard rate out of unemployment, h, as defined by Equation 3.

If this worker were to change behavior (lower their w^R) such that the average applied-for wage drops by some amount, how much faster would they find a job? A natural way to summarize the answer to this question is to consider the elasticity of the hazard rate, h, with respect to the average applied-for wage w^*:

$$\eta^h_{w^*} = \frac{dh}{dw^*}\frac{w^*}{h} \tag{12}$$

It turns out that under weak conditions, there is a simple formula for computing this elasticity. Over the two subsections below, we first derive and discuss this formula and then discuss how to implement the formula on our data.

A formula for the elasticity of job finding with respect to applied-for wages

Proposition. *Assume that w^R is inside the support of F and that F is continuous at w^R. Equations 3 and 5 then imply*

$$\eta^h_{w^*} = -\frac{1}{1-\frac{w^R}{w^*}} \tag{13}$$

Proof. When F is continuous at w^R and w^R is inside the support of F, Equation 5 implies that the average applied-for wage is a continuous, strictly increasing function of the reservation wage in a neighborhood around w^R. Accordingly, a continuous inverse function also exists that defines the reservation wage as function of the average applied-for wage. Denote this function $r(\cdot)$. From Equation 3, the hazard rate can be then written as $h = \lambda(1 - F(r(w^*))$.

Now consider the derivative of h wrt. w^*. It is defined as the limit of the following difference quotient as $\delta \to 0$ (if the limit exists):

$$D(\delta) = \frac{\lambda\left(1 - F\left(r(w^* + \delta)\right)\right) - \lambda\left(1 - F\left(r(w^*)\right)\right.}{(w^* + \delta) - w^*}$$

Plugging in twice from Equation 5 in the denominator yields:

$$D(\delta) = \frac{\lambda\left(1 - F\left(r(w^* + \delta)\right)\right) - \lambda\left(1 - F\left(r(w^*)\right)\right)}{\frac{1}{1-F(r(w^*+\delta))}\int_{r(w^*+\delta)}^{\infty} w\, dF(w) - \frac{1}{1-F(r(w^*))}\int_{r(w^*)}^{\infty} w\, dF(w)}$$

Algebra then yields:

$$D(\delta) = -\frac{\lambda\left(1 - F\left(r(w^*)\right)\right)}{w^* + \delta - \frac{\int_{r(w^*)}^{r(w^*+\delta)} w\, dF(w)}{F(r(w^*+\delta))-F(r(w^*))}}$$

Now consider the last term in the denominator. Bounding the integral from above and below yields:

$$r(w^*) \leq \frac{\int_{r(w^*)}^{r(w^*+\delta)} w\, dF(w)}{F(r(w^* + \delta)) - F(r(w^*))} \leq r(w^* + \delta)$$

Since r is continuous it follows that this term converges to $r(w^*)$ as $\delta \to 0$. It follows that the difference quotient also converges so that $\frac{dh}{dw^*}$ exists and can be written:

$$\frac{dh}{dw^*} = -\frac{\lambda\left(1 - F\left(r(w^*)\right)\right)}{w^* - r(w^*)}$$

Plugging this into Equation 12, using the definition of the hazard rate and recalling $r(w^*) = w^R$ then yields:

$$\eta_{w^*}^{h} = -\frac{w^*}{w^* - w^R} = -\frac{1}{1 - \frac{w^R}{w^*}}$$

□

Before proceeding, we make two remarks regarding this formula. First, we emphasize that the formula relies on the existence of an offer distribution and a reservation wage, on the definition of the hazard rate and on the definition of the average-applied-for wage. It does *not* rely on the reservation wage equation or on the form of the worker's Bellman equation. This makes the formula quite general as it does not rely on any assumptions about *how* workers choose their reservation wage. The formula is thus robust to a range of changes to the benchmark model, including letting the offer arrival rate depend on search effort, adding job destruction, on-the-job search, behavioral preferences or even non-optimizing agents that set their reservation wage using rules of thumb or other exogenous decision processes.[21]

Second, in terms of the underlying mechanics, the formula relies on basic properties of probability distributions and the expectation of truncated random variables (such as the average applied-for wage). The key is that when w^R and w^* are close, a given change in w^* corresponds to a larger change in w^R and thus to a larger change in the set of acceptable jobs. To see why, note that increases

[21] If search effort is assumed to affect the job finding rate, the hazard rate out of unemployment will of course depend both on the level of search effort and on the decision about which wages to target. $\eta_{w^*}^{h}$ can still be computed using Equation 13 however, and will show how the hazard rate change with applied-for wages when search effort is held fixed.

in w^R mean that workers stop applying for the lowest paying jobs in the existing distribution of applied-for (acceptable) jobs. This of course causes the mean of applied-for wages to increase. The magnitude of this increases depends on how close w^R is to w^* however. If wages in the lowest acceptable jobs are very close to the existing mean (w^R is close w^*) the new mean will be only slightly higher. Accordingly, to create some given change in w^* we need to increase w^R by a lot in this case. If instead w^R is far from w^*, we can create the same change by only changing w^R slightly.[22] This explains why the job finding rate changes relatively slowly with w^* when w^R is far from w^* and vice versa when w^R is close to w^*.

Empirical implementation of the formula

The formula in Equation 13 offers a simple empirical gauge of how the job finding hazard is affected by workers choosing to target different wage levels. To compute how a one percent decrease in the average applied-for wage affects job finding rates for a given worker, we simply need a measure of their average applied-for wage, w^*, and their reservation wage, w^R. Our longitudinal data on job applications offers a stratightforward way to obtain these. For a given worker, the average wage of the jobs they apply to can be used as an estimate of their w^*, while their lowest applied-for wage in the data can be used as an estimate of w^R.[23] Importantly, because this empirical implementation applies Equation 13 at the individual level it is valid under arbitrary (time-invariant) worker heterogeneity. This is important in light of the the large heterogeneity that we document in search behavior. [24]

Implementing Equation 13 on our data raises two practical considerations, both pertaining to the use of an individual's smallest applied-for wage as an estimate of their reservation wage. The first consideration is sensitivity to measurement error and outliers. If measurement error leads to even a single artificially low applied-for wage for some individuals this will cause us to erroneously infer a correspondingly low reservation wage for these individuals. To address this potential concern, we check the robustness of our results to winsorizing the bottom and top 1 percent of wages before

[22] To expand on this more formally, we can decompose the elasticity of interest as the product of two other elasticities (assuming they exist): the elasticity of the reservation wage with respect to the average applied for wage and the elasticity of the hazard rate with respect to the reservation wage, $\eta^h_{w^*} = \eta^h_{w^R} \cdot \eta^{w^R}_{w^*}$. The key point is that $\eta^{w^R}_{w^*}$ turns out to be proportional to $\frac{1}{1-\frac{w^R}{w^*}}$ because of the intuition given in the main text; when w^R and w^* are close, a given change in w^*translates to a larger change in w^R.

[23] More formally, under our assumptions (Equation 4 in particular), a worker's average applied-for wage in the data is an unbiased estimate of w^*, while their lowest applied-for wage in the data is a super-consistent estimator for w^R (converging to the truth at rate N instead of the usual $\sqrt{N}$), see e.g. Christensen and Kiefer (1991).

[24] An alternative empirical implementation of Equation 13 would be to use aggregate data on job outcomes (or applications) from all individuals to estimate the overall "mean-min" ratio of accepted (or applied-for) wages as in Hornstein et al. (2011). This approach assumes that workers are homogeneous and have the same reservation wage and average applied-for wage. Our results on variation in applied-for wages strongly reject this assumption; applied-for wages vary substantially across workers. Moreover, because much of the variation is unrelated to observable characteristics (see Section 4.3), it is also not possible to address this heterogeneity simply by conditioning on observables. In terms of the underlying theory, it might seem like the elasticity formula in Equation 13 should be closely related to the key results in Hornstein et al. (2011) because a version of the mean-min ratio appears in both. This is in fact not the case. The results in Hornstein et al. (2011) are based on cleverly exploiting the structure of the search model's Bellman equation, whereas our formula Equation 13 does not rely on the Bellman equation at all.

Table 5: The elasticity of the job finding hazard with respect to the average applied-for wage, sample median

	Raw	Winsorized	Trimmed	Individuals
At least 10 applications	5,641	5,926	7,659	84714
At least 50 applications	4,410	4,762	5,367	15414
At least 100 applications	4,011	4,527	4,764	2429

Notes: This table presents the median elasticity of job finding with respect to average applied-for wage in our sample (based on Equation 8). The three rows of the table correspond to different sample restrictions regarding the number of observed applications per unemployment spell. The three columns correspond to different approaches to outliers in wages. The first column simply uses the raw data, the second column, winsorizes applied-for wages at 1st and 99th percentile and the third column trims (drops) the largest and smallest applied-for wage for each individual in the data.

applying Equation 13.[25] We also consider a trimming procedure that simply drops the single highest and lowest applied-for wage for each individual.

The second consideration regards the number of applications observed for each individual. For individuals where we only observe very few applications, the smallest applied-for wage is unlikely to be a good estimate of their reservation wage.[26] To address this concern, we restrict the sample to require various minimum levels of observed applications. The downside to choosing a larger minimum number, of course, is that it both cuts our sample size and potentially tilts our sample towards a selected subset of UI recipients with longer spells.

In Table 5 we show the median elasticity of job finding with respect to average applied-for wage in our sample.[27] As discussed above, this is based on applying Equation 13 to each individual in our data. The three rows of the table correspond to different sample restrictions regarding the number of observed applications. while the first three columns of the table shows how results change when using raw wage data or when addressing outlier wages via winsorizing or trimming as described above. In line with expectations, we see that requiring more applications leads to a slightly lower median elasticity, while addressing outliers leads to higher elasticities. Across all implementations, however, we find that changes in applied-for wages are quantitatively important for job finding; the median elasticity ranges between 4.01 and 7.66.

[25] Obviously it is only winsorizing at the bottom that matters for our estimation of w^R. Winsorizing only at the bottom however, would mean that our estimate of w^* is asymmetrically affected and mechanically shifted up. We therefore take the internally more consistent approach of winsorizing symmetrically.

[26] To put this more formally, the smallest applied-for wage does converge fast to w^R as the number of observed applications grows (see footnote 23). In any finite sample of applications, however, it will be systematically biased towards w^* and in small samples, the bias can be substantial.

[27] We focus on the median instead of the mean to minimize the influence of individuals with very large $\frac{w^R}{w^*}$-ratios. Mean elasticities are similar in magnitude but slightly larger throughout.

Part II

Online Appendix

B Data construction and validity checks

B.1 Details on the Institutional Setting

To remain eligible for UI while unemployed, UI recipients have to be actively searching and applying for jobs. Additional requirements for maintaining eligibility are that the UI recipient accepts appropriate job offers and participate in activities (such as meetings and activation programs) at the municipal job centers and at the UI funds.

The administration and payout of UI in Denmark is carried out by the UI funds. This includes the administration of the job search documentation requirements in Joblog. During a UI recipient's first weeks of unemployment, the UI fund is legally required to instruct the UI recipient in the use of the Joblog system.[28] Over the subsequent unemployment spell, the fund is required to assess whether the UI recipient is complying with the documentation requirements necessary to maintain eligibility.

The Joblog system works as follows: To register applied-for jobs in the system, unemployed workers need to log in to the central online platform of the Danish public employment service *(Jobnet)*. This platform is also where UI recipients register for UI benefits at the start of their unemployment spell and where they book meetings with caseworkers at the municipal job centers. The website also serves as a job board with posted vacancies. Through the Jobnet platform, UI recipients then enter the Joblog system.

The law on UI requires that the UI fund always specifies a minimum amount of weekly or monthly applications that each individual needs to register. However, this amount should in principle be based on a specific assessment of the workers' education, work experience and competencies, as well as the demand for labor in the area that the worker needs to be available for. Note that the UI recipients also have to regularly attach the actual application as well to the *job-log* (i.e., the registered application in the Joblog system).[29]

Despite the lack of a formal universal threshold of registration requirements, the vast majority UI funds often post general guidelines of their expectations, and it is generally well-known that registering between 1.5 and 2 applications per week should be sufficient for recipients to fulfill eligibility requirements (see also footnote 28). For comparison, similar types of activity requirements

[28] The UI funds are incentivized to comply with the rules by the National Labor Market Authorities. If the National Labor Market Authorities decide that a UI fund has not administered according to the law (i.e., assessing eligibility and screening registered applications) and thus paid out "illegal" UI benefits, the UI fund risk losing the reimbursements of UI. The National Labor Market Authorities further use the Joblog data to create monthly scoring cards (or performance assessments) across different UI funds and municipalities on the share of UI recipients satisfying the UI eligibility requirements and logging at least 1.5 applications per week, see https://va.star.dk/.

[29] In our data, we know that around a third of job-logs contain an actual attachment, but we cannot see the actual content of the attachment. Further, for half of all the job-logs, there is also a link to a specific vacancy in the Jobnet portal.

exist in several other UI systems, see for example Marinescu (2017) and McVicar (2020).

Overall UI recipients face very clear incentives to register job applications in the Joblog system and to do so truthfully. As discussed above, UI recipients face sanctions if they fail to register the required number of applications or if they are caught registering fictitious applications.[30] These incentives are borne out in a very high level of registration activity in our data. In the raw data 96 percent of new UI recipients register at least one applied-for job during the UI spell. Among UI spells lasting at least 8 weeks - which we focus on - the number is even higher (see also Section 3.4).

B.2 Sample Selection and sample statistics

Our baseline sample is constructed from administrative data on UI payments and consists of all individuals of Danish origin entering new unemployment spells with full UI eligibility starting between September 2015 to September 2017.[31] We link all UI spells to data from a wide range of other administrative data sets maintained by Statistics Denmark (DST). These data sources are IDA, BFL and DREAM. IDA, the Integrated Database for Labor Market Research, is a matched employer-employee panel containing socioeconomic information on the entire Danish population. BFL, the Employment Statistics for Employees, contains data on jobs, paid hours of work and earnings for the universe of employed individuals. DREAM is an event-history data set created by the Ministry of Employment tracing the participation of individuals in public income support programs at a weekly level. All these data sets are available through servers at Statistics Denmark. From these data we extract demographic information, education and the full history of public benefit payments and employment, including information on occupation,[32] hours,[33] wages, industry[34] and firm identifiers for the employing firms.

We use the DREAM register database to identify new UI spells[35] and focus on spells with full

[30] Based on publicly available statistics on the sanctioning rate across different UI funds, the rate was estimated to be around 5 percent in the second quarter of 2021. Unfortunately, we are not able to separately identify the subset of sanctions which are related to missing search and documentation requirements. Therefore this number is an upper bound as it includes sanctions due to e.g. missing meetings at the job center or other forms of non-compliance. The size of sanctions related to failure to meet eligibility criteria typically ranges from a loss of benefits for two to three days to three weeks. In severe cases, benefits can be removed until new eligibility through employment has been established.

[31] The start and end sampling points reflect that we observe job-log entries from September 2015 and have labor market data available until September 2017. The initial panel that we construct is on the spell level, thus some individuals may enter multiple times.

[32] DISCO is the Danish equivalence of the standard international ISCO classifications, e.g. medical doctors have the minor code 221 which relates to the sub-major group of health professionals with code 22, which is part of the major group of professionals with number 2. Similar classifications apply for the industries. The first four digits of DISCO are identical to the international ISCO classification, details are here `https://www.dst.dk/da/Statistik/dokumentation/nomenklaturer/disco-08`

[33] Full-time positions are always defined as 37 hours per week in the Danish context. The exact number of hours involved in a part-time position varies in the Danish labor market. Sometimes anything below 37 hours can be categorized as part-time work. Part-time work with 20 hours per week is however quite common.

[34] We use the NACE Rev. 2 nomenclature to classify industries. Information in the aggregation of the NACE Rev. 2 nomenclature for industries can be found here: `https://www.dst.dk/klassifikationsbilag/8cf95f88-8153-43b5-a82a-fa89adf6f214` (pp. 463-477)

[35] Defined by observing at least 4 weeks of consecutive UI benefit payments without any UI payout in the 4 weeks prior to spell start.

UI eligibility. As special UI rules may apply to immigrants, we exclude them from the sample. For each UI recipient and unemployment spell, we then use a unique person identifier to identify all applied-for jobs that they have registered in the Joblog system during their unemployment spell.[36] To this baseline sample of spells with full UI entitlement we impose the following sample restrictions:

First, we impose the restriction that the UI spell lasts at least 8 weeks. We do this to remove individuals who are de facto making a job-to-job transition, but who are temporarily receiving UI while waiting for their new job to start. Since we do not consider spells with less than 8 weeks of duration the latest possible entry in our data is thus June 2017. Second, we discard information on applications submitted in the last 4 weeks of a spell. This gets rid of applications that UI recipients are making after successfully landing a job but before this job has actually started.[37] We impose the restriction on all spells to ensure that they have a similar observation period. Third, we require that the individual has registered at least one applied-for job during the observed unemployment spell so that some information on search behavior is available. Because of the high level of usage of the Joblog system, this only removes very few individuals. These three sample restrictions imply that only 434 individuals in our sample enter with multiple spells. To simplify the presentation of our sample, we therefore impose the fourth sample restriction to only include the first unemployment spell for each individual. Finally, as the last step we focus on search behavior for these spells during the first 12 months. See Section 3.2 for a discussion of this last step.

Table 6 shows the effect of these additional sample restrictions in terms of the reduction in sample size. Given these restrictions, we arrive at our final analysis sample consisting of 127.695 individuals and 4 million applications. In Figure 3 we show histograms of the number of logged job applications at the monthly or weekly level of aggregation. Clearly the distribution of the number of logged applications centers around the semi-official registration incentives set forward by the UI funds, see Sections 3.4 and 3.1.

In Figure 4 we show the evolution in the number of logs over time in unemployment. The figure reports estimated duration fixed effects from regressions with and without individual fixed effects (Equations 6 and 7 respectively). The raw time profile suggests that the average number of logged applications is fairly stable throughout time in unemployment after the first month or so.[38]

[36] Note that the Joblog system also allows the UI recipients to register additional information beyond the required information. This additional information includes registering jobs that they plan to apply for in the future and registering other activities such as participating in a job interview. Since UI recipients are not required to use most of these features, fewer UI recipients register these activities. In our analysis, we only use data on jobs that UI recipients report having applied for.

[37] Many jobs do not start right away which implies that UI recipients typically continue receiving UI for some weeks after they have accepted a new job. In the raw data, we see a clear drop in the number of applications that people register in Joblog about one month before they enter employment, likely reflecting that the individuals have already accepted their new job at this point in time and are simply waiting for it to start. As a consequence of the wording of the Danish UI rules during our sample frame, however, such individuals were in principle required to both apply for and register applications in Joblog, even though they had a new job lined up with a known start date. They obviously face a very peculiar set of incentives in their application decisions. By dropping applied-for jobs in the last 4 weeks before a new job starts, we get rid of these applications.

[38] Figure 4 shows an initial phasing-in period where the number of applications in the first month is about one third lower than in the following month. Note that, especially, in the early periods of unemployment there are good reasons to expect that the logging requirements are not as binding. At the beginning of a new unemployment spell,

Table 6: Sample selection

	Individuals	Spells	Job-logs
Initial sample with full UI entitlement	170,890	172,110	5,232,416
At least 2 months (8 weeks) in UI	130,086	130,554	4,975,547
Censoring last 4 weeks of applications	130,086	130,554	4,316,519
Min. 1 logged application	127,695	128,129	4,316,519
First spell	127,695	127,695	4,307,886
Remove logs after 12 months	127,695	127,695	4,013,475
Final sample	127,695	127,695	4,013,475

Notes: The Table shows the amount of individuals and respective unemployment spells as well as number of job logs for various stages in our sample selection process. The last row indicates the final sample used for the main analysis of this paper.

However, when we look at the estimates from the regression including person fixed effects, we find that registered applications decrease over time. This difference in estimates suggests that individuals with longer unemployment durations also register more applications than individuals who find jobs earlier on, which is similar to the result of Faberman and Kudlyak (2019). Nevertheless, it is difficult to say whether these dynamic selection patterns reflect true differences in search effort across individuals, or whether it is related to individuals changing the share of their applications they register over time.

In Table 8 we describe our sample in terms of average applied-for job characteristics across months in our final sample and across the different unemployment duration groups similar to Table 4. As explained below there are some applications which we cannot link to firms in the registers, or where information on previous occupation or applied-for occupations is not available. We treat these observations as missing when constructing the monthly averages. In Table 8 we report summary statistics on this dimension also by e.g. the share of observations (months) where we have no information on typical wages.

B.3 Matching algorithm

We now provide additional details on how we process the job-log entries to obtain links to occupation codes and firm identifiers which enables linking to the registers available at Statistics Denmark. Before matching reported job titles and firm names to official classifications and registers, we perform extensive cleaning of these entries. In this step, we streamline the notation between source and target files and correct basic spelling mistakes.

As a first step, we use the self-reported job titles and link these to the official Danish occupational codes (DISCO, see footnote 32). We exploit that many of the self-reported job titles have the actual

the unemployed are usually subject to a 'phasing-in' period in which individuals slowly get introduced to joblog and other components of the UI system (keep in mind that the job log system is still relatively new in our sample and that a very large part of individuals in our sample have limited previous unemployment history (see Table 4)). Hence, over the first weeks, more time is spent on meetings that prepare job search such as composing a proper CV with the help of caseworkers etc.

Figure 3: Histograms over number of applications

(a) Monthly level

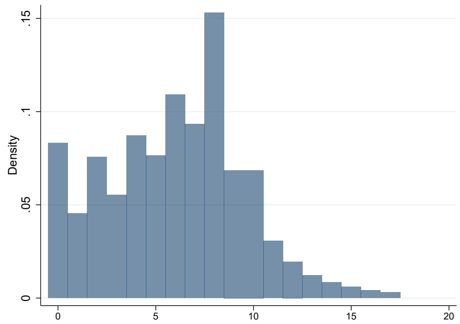

(b) Monthly level - excl. first month on UI

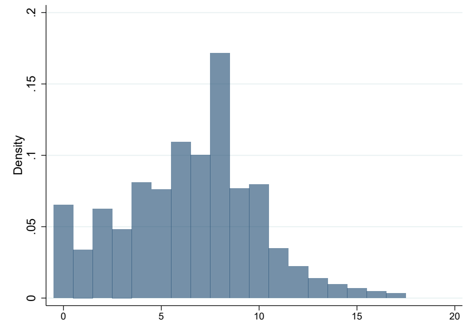

(c) Weekly level

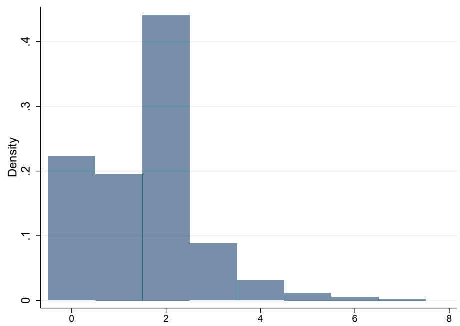

(d) Weekly level - conditional on logging

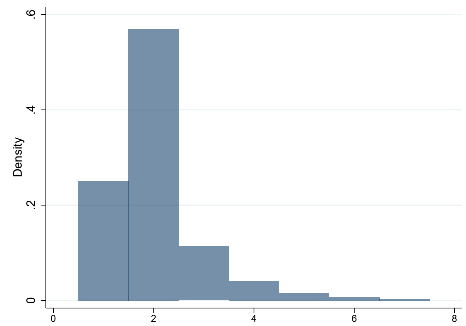

Note: The figure shows histograms of the distribution of the number of logged applications at the monthly or weekly level of aggregation. We further show a version where we exclude submitted applications in the first month of UI and a version where we only include weeks where the individual logged at least one application. Due to Statistics Denmark's confidentiality rules, we have excluded observations above the 99th percentile in all figures (18 applications at the monthly level).

name of the occupation as a part of the self-reported title. Thus, we identify occurrences of the DISCO occupations in the reported job titles. We only consider 1:1 matches (43.4 percent), i.e., if a certain job title links to several occupations, we do not treat it as a match. For remaining unmatched entries we manually match some job titles to occupations as many job titles use acronyms that do not match the standard classification (i.e., 'social og sundhedshjælper', Danish for social and health-care workers, are most often reported as 'sosu-hjælper'). This adds about 27.2 percent to the matches. Finally, we also use some fuzzy matching techniques on the remaining unmatched observations to circumvent misspellings in the job-log job titles, adding the manual titles from the step before. We rank the potential matches along several scoring functions and only pick consistently high ranked matches. For this, we use compget, speedist and soundex routines from SAS as well as sub-string occurrences which adds 10.9 percent. Overall, we can thus map 81.5 percent of the applications to a DISCO group.

In the second step, we link the reported firm information to official firm identifiers. With the mandatory reporting of firm name, zip code and city, we develop a matching procedure which

Figure 4: Number of monthly logged applications

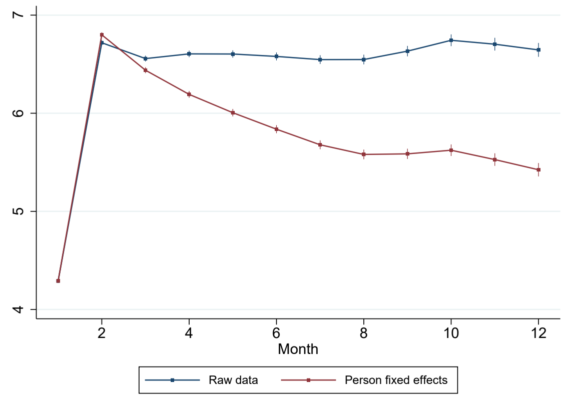

Note: This figure shows the estimates of the duration fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, estimated from Equation 6 (blue line) and Equation 7 (red line). Note that in the red line we have normalized τ_1 to equal the mean of the outcome variable in the first month of unemployment. The outcome variable is the number of logged applications in a given month. Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

matches this information to the official firm registers recording all Danish firms (CVR-register). We can then use these links to identify firms in the registers at Statistics Denmark. Our matching procedure on firms starts with "perfect matches", using both firm names and zip codes. Here we have a 1:1 match for 66.3 percent of the applications in the Joblog data. We further add the sub-string matches (i.e., where we have a "perfect match" for a subset of the firm name string), and if several potential matches exist, we choose the one which is spatially closest to the reported firm address. This adds 13.9 percent to the matches. To link applications which we cannot match exactly on firm names, we employ a fuzzy matching procedure using the matchit command in STATA to identify the 50 closest matches. We then test these 50 potential matches using several scoring functions besides the one obtained from matchit. For each of the scores (5 in total), we calculate the ranking of the 50 potential matches and identify the "correct'' match as the match which receives the best ranking on average (the scores we use are Bi-gram Similscore, Token, TokenSound from matchit and the compget and speedist functions in SAS). This adds a further 6.2 percent to the matches, so we end up with an overall firm match of 86.4 percent.

B.4 Information on applied-for wages and firm type

Our data do not contain wages associated with the applied-for job as these are not commonly reported for vacancies in Denmark. Instead we construct two different proxies for the wage associated with a given applied-for job leveraging that we have access to the rich Danish register data.

First, we use estimates from an AKM model (Abowd et al. (1999)) and use the implied firm fixed effects as a measure of the firm type. We take advantage of the rich administrative data on the whole Danish working population (the BFL data, see Section B.2) covering all salaries in Denmark to construct a matched employer-employee panel from 2008 to 2015 with 290,108 (connected) firms.

Table 7: Summary stats on job search

	Full sample		
	Mean	Std	Observations
Avg. typical wage	5.18	0.13	575,504
Avg. commute time	44.72	35.03	591,333
Share of full-time applications	0.89	0.22	597,268
Occupation unrelated to prev. job	0.43	0.4	497,890
Industry unrelated to prev. job	0.52	0.38	546,006
Scope of search index	0.08	0.98	466,878
Publicly posted vacancy	0.74	0.33	597,268
Network	0.06	0.16	597,268
Uninvited applications	0.19	0.29	597,268

Note: This table reports means and standard deviations on measures of monthly job search behavior in our data. See Section 3.3 for further details on these outcomes. Note that we have aggregated the data to the monthly level, i.e., an observation is the job search of individual i in month m.

We include year-month fixed effects in the AKM wage regression to absorb any aggregate time trends. Post-estimation we remove (worker-weighted) industry-specific means from the estimated firm effects.

Second, for each application in our data we estimate the typical wage for this position based on detailed observable characteristics of the job and the re-employment wages individuals are paid upon entering a similar type of job. In constructing this measure, we use occupation and industry codes of different levels of aggregation as well as AKM firm fixed effects. Further, we also include various interactions between all of these measures and use a LASSO regression for model selection (Belloni et al., 2014). Specifically, we consider a linear regression with log wages as the outcome variable and a very large number of potential explanatory variables based on the available job characteristics in our data. We then use LASSO estimation to select the subset of these variables that most efficiently trade off predictive power in-sample against the risk of over-fitting. We rely on the Rigorous-LASSO approach of Belloni et al. (2016) to choose the regularization parameters for the LASSO estimation. We allow for clustered disturbances at the individual level in the estimation.

The estimation was conducted using the LASSOPACK implementation of Ahrens et al. (2019a,b). Out of the 10,407 baseline explanatory variables,[39] the Rigorous-LASSO selects 233 variables. As the final step, we run a standard OLS regression with log wage as the outcome variable and these 233 variables as explanatory variables (so-called Post-LASSO OLS) to arrive at our final prediction model. Further details on the estimation procedure for typical wages are available in Fluchtmann et al. (2019). These typical wages are then a measure of the wage a person would typically be paid in the applied-for job.

[39]As the baseline set of explanatory variables we include dummies for the industry and occupation of the job at both the 1-, 2- and 3- digit levels, a dummy for whether we were able to obtain an estimated AKM fixed effect for the employing firm, and the within-industry-demeaned AKM firm fixed effect when this is available. In addition we include all pairwise interactions between these variables.

Table 8: Summary statistics across monthly averages - Measures of applied-for job characteristics

	Unemployment duration					
	Total	**< 4 Months**	**4-6 Months**	**7-10 Months**	**> 10 Months**	**Right Censored**†
Typical Wage (log)	5.187	5.186	5.181	5.181	5.185	5.201
	(0.120)	*(0.123)*	*(0.117)*	*(0.115)*	*(0.115)*	*(0.121)*
AKM firm fixed effect	0.00930	0.00741	0.00627	0.0111	0.0116	0.0148
	(0.451)	*(0.559)*	*(0.402)*	*(0.321)*	*(0.288)*	*(0.428)*
Fulltime	0.885	0.887	0.886	0.881	0.888	0.881
	(0.192)	*(0.206)*	*(0.185)*	*(0.178)*	*(0.169)*	*(0.192)*
Unrelated Occupation (0/1)	0.412	0.382	0.410	0.428	0.460	0.438
	(0.370)	*(0.383)*	*(0.359)*	*(0.351)*	*(0.353)*	*(0.376)*
Unrelated Industry (0/1)	0.492	0.466	0.495	0.513	0.538	0.499
	(0.340)	*(0.369)*	*(0.327)*	*(0.306)*	*(0.293)*	*(0.337)*
No information on typical wages	0.119	0.155	0.125	0.112	0.0785	0.0690
	(0.185)	*(0.222)*	*(0.174)*	*(0.154)*	*(0.129)*	*(0.146)*
No information on firm effects	0.121	0.159	0.127	0.112	0.0780	0.0674
	(0.186)	*(0.223)*	*(0.175)*	*(0.154)*	*(0.129)*	*(0.144)*
No information on hours	0.0849	0.110	0.0881	0.0817	0.0583	0.0473
	(0.161)	*(0.195)*	*(0.149)*	*(0.137)*	*(0.115)*	*(0.122)*
No information on commuting time	0.0950	0.124	0.0979	0.0900	0.0646	0.0542
	(0.170)	*(0.207)*	*(0.156)*	*(0.141)*	*(0.118)*	*(0.132)*
No information on unrelated occupations	0.110	0.140	0.115	0.105	0.0773	0.0653
	(0.181)	*(0.214)*	*(0.169)*	*(0.155)*	*(0.138)*	*(0.143)*
No information on unrelated industries	0.120	0.159	0.127	0.112	0.0777	0.0671
	(0.186)	*(0.223)*	*(0.175)*	*(0.154)*	*(0.129)*	*(0.143)*
Avg. number of logs per individual	31.52	11.78	29.10	52.32	81.78	27.26
	(28.15)	*(7.272)*	*(14.52)*	*(21.96)*	*(27.58)*	*(21.67)*
Avg. number of monthly logs	5.80	4.79	5.96	6.39	6.99	6.51
	(2.78)	(2.68)	(2.84)	(2.50)	(2.33)	(2.71)
Number of monthly logs	4,013,475	524,034	859,313	747,396	1,285,839	596,893

Notes: The table displays the average characteristics of the search portfolio over the first 12 months of unemployment. Averages are calculated across the monthly averages/measures of applied-for job characteristics. Standard deviations are reported in parentheses. † Censored spells are defined as spells which are censored with duration less than 11 months

B.5 Building commuting, occupation and industry relatedness measures

Using the firm's address reported in the Joblog data and the individual's residence municipality, we determine commuting times. The commuting times are based on distance measures from Google maps API, these data come from Harmon (2013).

Our analysis also focuses on the occupation and industry of the applied-for jobs. Here we specifically focus on the proximity (or similarity) of applied-for jobs' industry and occupation relative to the workers' previous jobs. The motivation is that jobs which require a very different set of skills to that of the job seekers' previous jobs are likely less attractive to apply for. For occupations, we base our measure of relatedness on the latest version of the O*Net Related Occupations Matrix. These data contain, for each occupation, the top 10 related occupations in terms of skills and experience (Allen et al., 2012). We map this matrix to our 3-digit DISCO codes and use this to define the set of related occupations.[40] Based on our measure of related occupations we then say that an occupation is *unrelated* to the job seekers' past occupation if it is not in the set of related occupations. To test the robustness of our findings based on this measure, we also construct an alternative measure as explained in Appendix C.3.1.

To get a similar measure for skill-relatedness across industries, we use data from Neffke and Henning (2013). These data contain skill-relatedness estimates across NACE Rev. 2 industries based on labor flows among industries in the Swedish economy. We select the top 10 most related 3-digit industries to resemble the occupation measure and define this set as the group of related industries. Based on these measures, we say that an occupation or industry is *unrelated* to the job seekers' past occupation or industry if it is not among the 10 most related occupations or industries.

As a final step we standardize[41] our three measures (commuting time, industry and occupation relatedness) and combine them into a scope of search index. This constructed scope of search index is finally standardized by the mean and standard deviation from the first month of search, such that scope of search (Figure 1c) evolve from mean 0, and changes are evaluated in standard deviations of the variation across all individuals in their first month of search.

B.6 Coverage and validity of Joblog data

In this section we complement the discussion in Section 3.4 with additional details regarding the coverage and validity of Joblog data. First, we discuss evidence from a survey regarding the coverage of our Joblog data. Second, we examine how often we are able to trace a new hire back to a job application that is contained in our data and discuss the implications. Third, we look at how the Joblog data on applied-for jobs relate to the actual hiring outcomes we observe in the data. Fourth,

[40]The O*Net Related Occupations Matrix is based on US data. In this classification, plumbers are e.g. coded to be highly related to the occupation of heating and air condition mechanics. The matrix defines related occupations in terms of tasks and requirements using a classification of occupations that we can map to the Danish 3-digit DISCO codes after appropriate translation. Some of these codes are more detailed than the DISCO codes. Sometimes, we therefore get over 10 related occupations for a single 3-digit DISCO code.

[41]Specifically, we standardize commuting time, industry and occupation measures so the variables have mean 0 and standard deviation 1 in month 1.

Table 9: Survey question "Which of these statements best describes your use of Joblog?"

Answer:	**Share of respondents:**
Fulfill requirements, often applied to more jobs	37%
Fulfill requirements, rarely applied to more jobs	21%
Always register all applied-for jobs	41%
Never register applications	1%
Number of respondents	1236

Notes: The table shows answers to the question "Which of these statements best describes your use of Joblog?" based on the survey of UI recipients conducted in Mahlstedt et al. (2019).

as a final robustness check, we repeat the key results from our main analysis (Figure 1) only for those individuals in our sample who are eventually hired and where we can find a past application from the same individual to the firm.

Note that some of the analysis below is adapted from Fluchtmann et al. (2019) in which we use the same data sources to study gender differences in job search.

B.6.1 Survey

Tables 9 and 10 present results from a survey conducted among Danish UI recipients by Altmann et al. (2019) in March 2018.[42] Table 9 reports survey answers about how individuals log applications in Joblog. 41 percent of respondents report that they always log all the jobs they have applied for in Joblog, regardless of whether they have fulfilled the logging requirements. An additional 21 percent report that they only log applications up to the point where they have satisfied their logging requirements, but that they rarely apply for more jobs than what is required. Putting these together suggests that Joblog has close to full coverage for 63 percent of respondents. For the remaining 37 percent, however, the survey responses suggest that the Joblog data often miss some job applications that they have made beyond the required number.

To get a sense of how many applications may be missed by the Joblog data, Table 10 presents survey responses about the total number of job applications sent the past month and the number of job applications sent that were not registered in Joblog. In addition, the table also shows the actual number of registered applications (i.e. job-logs) made by the survey respondents in the month before the survey. This was computed by linking survey responses with the actual Joblog data. On average, survey respondents report applying for 11.5 jobs in total over the past month. Of those jobs, survey respondents on average say they failed to register 2.4 jobs in Joblog. This suggests that Joblog covers 80 percent of actual applications. The bottom of the table instead shows that the average number of jobs respondents actually registered in Joblog was 8.0. Relative to the total number of reported applications, this suggests that respondents on average failed to register 3.5 applications, implying that Joblog on average covers 69 percent of all applications.[43]

[42]We thank the authors for making these data available.

[43]This difference could reflect imperfect recall among survey respondents or could relate to measurement error from

Table 10: Self-reported and registered applications in the previous month

	Mean
Survey answers	
# of applied-for jobs	11.5
# of applied-for jobs not registered	2.4
Joblog data	
# of applied-for jobs	8.0

Notes: The top part of the table shows the reported number of job applications sent over the last month and the reported number of these jobs applications that were not registered in Joblog based on the survey of UI recipients conducted in Mahlstedt et al. (2019). The bottom part of the table shows the actual number of jobs registered in Joblog by the survey respondents in the month prior to the survey.

B.6.2 Successful applications

We also examine how often we are able to trace a new hire back to a job application that is contained in our data. For 47 percent of the new hires, we can identify a previous (logged) application by the UI recipient to the firm in question. This is informative about the representativeness of the data. To see why, assume that the Joblog data cover a share r of all applications and that the share of applications that we successfully match to firms in our data matching procedure is s. In this case, our data will contain firm information for a share $s \cdot r$ of all applied-for jobs.

Next assume that the fraction of jobs that stems from a job application is j. If the applied-for jobs in our data are a representative subset of all applied-for jobs, the share of new hires that we should be able to trace back to an application, t, should then be:

$$t = j \cdot r \cdot s$$

Based on independent survey data from Table 10, we estimated that the raw Joblog data contain between 69 and 80 percent of all applied-for jobs, i.e. r is between 0.69 and 0.80. Furthermore, as described in Section B.3, $s = 0.86$ in our data matching procedure. Finally, to provide bounds on the share of hires that stem from a job application, j, we rely on Statistics Denmark's official survey *Arbejdskraftsundersøgelsen* on how unemployed Danes report landing their first job out of unemployment (Engmann and Weiskopf (2019)).

In these data, 11 percent of respondents report landing their job in a way that is very unlikely to have involved the worker applying for the job (the job resulted from work at a temp agency,

the timing of registered jobs and/or the precise interpretation of the survey question. Registering applied-for jobs in Joblog can be done retroactively, so the interpretation of the survey question could either refer to the date at which applications were sent, or to the date at which the application was entered into the Joblog system. Additionally, the fact that UI recipients are able to register other activities besides formal job applications introduces some ambiguity about the interpretation of the survey question (if for example UI recipients have registered that they reached out to a friend about a specific job).

they got the job via their educational institution as an internship or the job seekers themselves advertised publicly), while 58 percent of respondents report landing their job in a way that almost surely involved making a formal job application (they themselves applied to a posted position, they applied to a firm with no posted positions, or they were directed to the job by the employment agency or other authorities).[44] To arrive at an estimate for the fraction of hires that stems from workers applying for the job, we simply assume that half of the remaining jobs involved a job application. This implies that about 73 percent of new hires out of unemployment involve the worker applying for the job at some point so that $j = 0.73$.[45]

Plugging in these values, we see that if the applied-for jobs in our data are representative, the share of new hires that we should be able to match, t, should be between $0.73 \cdot 0.69 \cdot 0.86 = 0.43$ and $0.73 \cdot 0.80 \cdot 0.86 = 0.50$. As noted, we in fact have $t = 0.47$ in our data, consistent with the data containing a representative subset of all applied-for jobs.

B.6.3 Joblogs and hiring outcomes

For the sub-sample of UI recipients who eventually find a job, we analyze how the Joblog data on applied-for jobs relate to the actual hiring outcomes we observe in the data. If differences in applied-for job characteristics across individuals also correlate with differences in the type of employment the unemployed enter, this is a strong signal that the Joblog data accurately capture actual job search behavior.

We benchmark the predictive value of the Joblog data against a known strong predictor of job outcomes: the characteristics of the UI recipient's previous job. In Table 11 we show how the characteristics of applied-for jobs and prior job characteristics predict respectively the industry, the occupation, the firm wage level (AKM firm fixed effect) or the wage level of a UI recipient's new job. Each column of the table corresponds to a different prediction model estimated on our analysis sample.

When predicting the industry of the new job, we use a simple multinomial logit model that includes either dummies for the industry of the previous job, the share of applications going to each industry or both sets of variables. Similarly, when predicting the occupation of the new job we use a multinomial logit model that includes either dummies for the occupation of the previous job or the share of applications going to each occupation. We observe some industries and occupations with very few new hires in our sample. Therefore we drop one small industry and two small occupations.[46]

[44] The remaining respondents report landing their jobs through channels that may have involved applying for the job application, but may also have involved receiving a job offer more directly. This includes finding the job through an acquaintance or finding a job after having been contacted by the firm.

[45] Alternatively, we could use 0.58 as a lower bound on r and use 0.89 as an upper bound. Plugging into the formulae above, we then see that if the applied-for jobs in our data are representative, the share of new hires that we should be able to match, t, should be between $0.58 \cdot 0.69 \cdot 0.86 = 0.34$ and $0.89 \cdot 0.80 \cdot 0.86 = 0.61$.

[46] Specifically, we exclude individuals from the sample who find a job in the smallest industry (less than 1.3 percent of new jobs) or the two smallest occupations (1.2 percent of new hires or less), respectively, as well as individuals whose prior job was in this industry or these occupations. Results are almost identical if these observations are included, however; we see indications that the likelihood function becomes ill-behaved in some specifications in this case, reflecting that some observations are predicted nearly perfectly.

When predicting the firm wage level (typical wage) of the new job, we use a simple linear regression that includes either the firm wage level (typical wage) of the previous job or includes the mean of the firm wage level (typical wage) across all applied-for jobs. For the linear regression models, we measure the predictive power simply using the regression R^2. For the multinomial logit models, we use McFadden's *pseudo*-R^2.

Looking across Table 11, we see that models that predict job outcomes only using data on applied-for jobs perform quite similarly to models that instead use prior job characteristics. The data on applied-for jobs do worse than prior job characteristics when predicting the firm wage level (columns (7) and (8)), but only slightly worse when predicting the occupation of the new job (column (4) vs (5)). At the same time, data on applied-for jobs perform almost just as well as prior job characteristics when predicting the typical wages on new jobs (columns (10) and (11)) and do better than prior job characteristics when predicting the industry of the new job (column (1) vs. (2)).

For models that include both prior job characteristics and data on applied-for jobs (columns (3), (6), (9) and (12)), we see that the characteristics of applied-for jobs remain highly predictive even after prior job characteristics have been conditioned on; adding the applied-for job variables alongside prior job characteristics always leads to sizable increases in the (*pseudo*-)R^2 relative to models that only use prior job characteristics. Moreover, the applied-for job variables are always highly statistically significant in the combined models. Overall, we conclude that the Joblog data are highly predictive of later job outcomes.

We also focus on the measured changes in job search behavior over time. Since these changes are at the heart of our analysis, they must reflect meaningful changes in behavior rather than some form of dynamic measurement error or changes in selective reporting of applied-for jobs over time. As a check on this, we examine whether application data from the end of the unemployment spell can predict later job outcomes, also *after* we first condition on application data from the early months of the spell. Specifically, we focus on the subset of our analysis sample that remains unemployed for at least 3 months but eventually find a job.

Table 12 shows the results of this type of exercise. In Column (1), we regress the typical wage of the new job on the average applied-for wage of the last two months while controlling for the average applied-for wage in all months prior to this. In Columns (2)-(11) we instead consider the industry of the new job. Here we regress a dummy variable for ending up in a particular 1-digit industry on the share of applications going to that industry over the last two months, while controlling for the share of applications going to the industry over all months prior to this. We see that job search behavior over the last two months of a spell is a strong predictor of the type of job that a UI recipient finds, even after conditioning on past job search behavior. For wages, a 1 percent increase in the average applied-for wage of the last two months is associated with a 0.29 percent increase in the wage paid in the new job. For industries, we see that a one percentage point increase in the share of applications going to a particular industry over the last two months is associated with an increase in the probability of ending up in this industry of between 0.19 and 0.30 percentage points.

Table 11: Predicting job outcomes from application data vs. prior job characteristics

Job Outcome:	Industry (1-digit)			Occupation (1-digit)			Firm wage level			Typical wage		
Model:	Multinomial logit			Multinomial logit			Linear regression			Linear regression		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Explanatory variables:												
Characteristics of previous job	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Characteristics of applied-for jobs	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes
# of parameters	80	80	152	63	77	133	2	2	3	2	2	3
(pseudo-)R-squared	0.287	0.237	0.356	0.329	0.389	0.492	0.048	0.063	0.094	0.256	0.272	0.377
p-value, test of excluding applied-for job variables	< 0.01		< 0.01	< 0.01		< 0.01	< 0.01		< 0.01	< 0.01		< 0.01

Notes: Columns (1)-(3) correspond to multinomial logit models for the 1-digit industry of the UI recipient's new job. Explanatory variables in these models are dummies for the 1-digit industry of the previous job or the share of job applications sent to jobs in each 1-digit industry. Columns (4)-(6) correspond to multinomial logit models for the 1-digit occupation of the UI recipient's new job. Explanatory variables in these models are dummies for the 1-digit occupation of the previous job or the share of job applications sent to jobs in each 1-digit occupation. Columns (7)-(9) correspond to linear regressions where the outcome variable is the industry-demeaned firm fixed effect for the UI recipient's new job. Explanatory variables in these models are the industry-demeaned firm fixed effect for the UI recipient's previous job or the average industry-demeaned firm fixed effect across all the applied-for jobs. Columns (10)-(12) correspond to linear regressions where the outcome variable is the typical wage of the UI recipient's new job. Explanatory variables in these models are the typical wage of the UI recipient's previous job or the average typical wage across all the applied-for jobs. The table reports the R^2 for the linear regression models. For the multinomial logit models, the table reports the McFadden's pseudo-R^2. The last row of the table shows the p-value for testing the exclusion of all explanatory variables pertaining to applied-for jobs.

These results confirm that measured changes in application behavior in our data are meaningful; individuals who show different changes in behavior over time also tend to have very different job outcomes. In other words, two job seekers who have the same application behavior in the first part of their UI spell but who later diverge in their search behavior also face very different job outcomes. This confirms that changes in application behavior over time in our data capture meaningful changes in behavior.

Finally, we note that the coefficient on search behavior in previous months in Table 12 shows that past search behavior also predicts later job outcomes. This is unsurprising. For some new jobs the gap between applying for the job and starting can be larger than 2 months, implying a direct link between past application behavior and later outcomes. Additionally, individuals who apply heavily to a particular type of job early in the spell are likely to stand out on other dimensions that make them likely to end up in this type of job also at a later date.

B.6.4 Robustness check: Requiring a successful link between application and new hire

In this subsection, we redo key parts of our main analysis only for those individuals in our sample who eventually find a job. We distinguish spells by whether we can link the new hire to a past application from the same individual to the firm (we call this application the successful application, see also Section B.6.2) or not. The motivation is that if some individuals are less diligent or truthful in registering applied-for jobs in Joblog, then these individuals should be (partially) removed by considering only individuals with a successful link.

Figure 5 estimates our event-study specification (Equation 7) across our four main outcomes as in Figure 1: the typical wage, the share of full-time jobs , scope of search and the share of jobs found via vacancy posting. Reassuringly, we find very similar results across samples, except for posted vacancies where we tend to see a slightly slower increase over time.

Table 12: Job outcomes and changes in applied-for job characteristics

	Log typical wage in new job (1)	Agricult., forestry and fish. (2)	Manufact., materials and util. (3)	Cons-truction (4)	Retail and transport. (5)	Information and communic. (6)	Finance and insurance (7)	Real estate and rental (8)	Business services (9)	Public adm., education and health (10)	Culture and leisure (11)
Avg. log wage of applied-for jobs, last 2 months	0.286*** (0.007)										
Avg. log wage of applied-for jobs, all other months	0.319*** (0.007)										
Share of applied-for jobs in this industry, last 2 months		0.297*** (0.006)	0.241*** (0.005)	0.292*** (0.005)	0.208*** (0.005)	0.265*** (0.005)	0.213*** (0.004)	0.186*** (0.005)	0.207*** (0.007)	0.272*** (0.006)	0.215*** (0.005)
Share of applied-for jobs in this industry, all other months		0.367*** (0.006)	0.254*** (0.005)	0.335*** (0.005)	0.206*** (0.005)	0.236*** (0.005)	0.206*** (0.004)	0.168*** (0.005)	0.205*** (0.007)	0.303*** (0.006)	0.234*** (0.006)
Observations	55151	84376	84376	84376	84376	84376	84376	84376	84376	84376	84376

Notes: Each column of the table shows OLS estimates from a different regression, estimated on the subset of the analysis sample that end their UI spell by finding a new job. In Column (1) the outcome variable is the log typical wage for the new job, and the right-hand side variables are the average log wage of applied-for jobs over the last 2 months of the UI spell and the average log wage of applied-for jobs over all other months of the spell. In each of the columns (2)-(11) the outcome variable is an indicator for whether the new job is in a different 1-digit industry, and the right-hand side variables are the share of applied-for jobs in this industry over the last 2 months of the spell and the share of applied-for jobs in this industry over all other months. Variation in observation numbers is caused by missing wage information in new jobs. Standard errors clustered at the individual level are in parentheses. *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

Figure 5: Results only using matched new hires

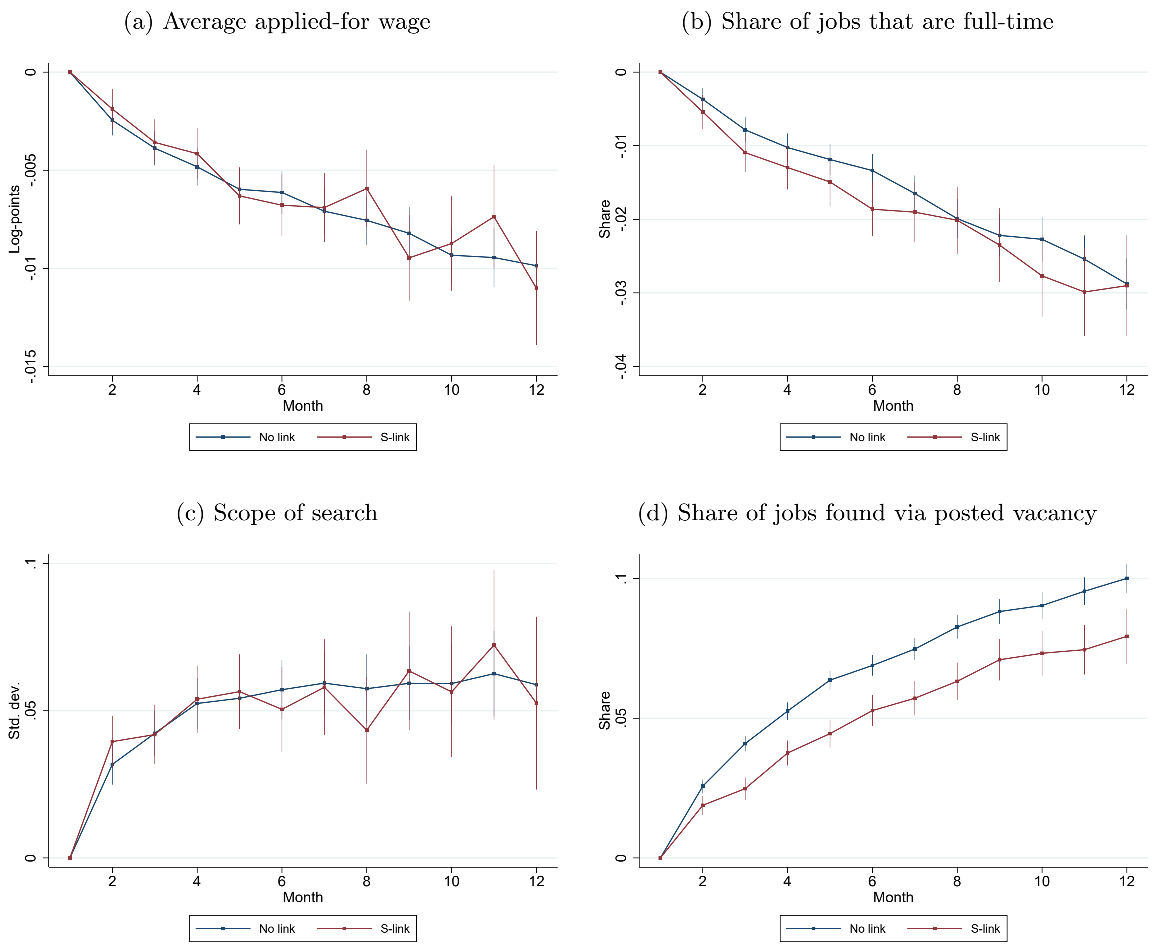

Note: The red (blue) line shows the estimates of the duration-month fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, estimated from Equation 7 for new hires where we can (cannot) link to a past application from the same individual to the firm. Note that for both lines we have normalized τ_1 to 0 in the first month of unemployment. The outcome variables are the average typical wage of applied-for jobs, share of applications to full-time jobs, a measure of broadness of search and the share of applications found via posted vacancies. See Section 3.3 for further details on these outcomes. Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

C Additional results

C.1 Firm type measures

In Figure 6 we show that the average applied-for firm type is decreasing throughout the unemployment spell, further supporting the notion that the unemployed workers gradually expand on the set of acceptable firms. We show the results for three alternative measures of the firm type. First, in Figure 6a the measure of a firm type is the average applied-for firm fixed effects from an AKM wage regression (see Section B.4). Second, in Figure 6b the measure is the average applied-for firm wage rank, which is determined by simply ranking firms by their average wage level across all employment spells in Denmark from 2008 to 2015 in the BFL data (see Section B.2). Third, in Figure 6c the measure is obtained by ranking firms by their value-added per worker when available. Value-added per worker is here determined by subtracting firm purchases from firm sales and dividing through by firm level employment (in full-time equivalents). The measures of firm sales and firm purchases are obtained from the FIKS database at Statistics Denmark,[47] and the measure of firm level employment is calculated based on the BFL data.

C.2 Dynamics in the distribution of applied-for wages

In this subsection we analyze how the decline in average applied-for wages throughout the unemployment spell (as documented in Figure 1) materializes through changes in the within-month distribution of applied-for wages. The fact that average applied-for wages are gradually declining over time is in line with theoretical predictions from standard search models. However, different theoretical frameworks make different predictions about how the gradual drop in applied-for wages should come about.

In a random search mdoel where a worker decision is summarized by a reservation wage decision the effect of a decreasing reservation wage would imply that job seekers now have a positive probability of applying to jobs/accepting wage offers below the initial truncation point (reservation wage).[48] However, job seekers do not change their decision about whether to accept/apply for job offers way above the reservation wage. For a large enough set of applications we should therefore expect to see very large wage offers throughout the unemployment spell, and we would expect to see less substantial changes in the set of the highest applied-for wages.[49]

[47]The FIKS database is constructed based on the mandatory VAT reports which are submitted from the universe of (VAT eligible) Danish Firms, for further information see: https://www.dst.dk/da/TilSalg/Forskningsservice/Dokumentation/hoejkvalitetsvariable/firmaernes-koeb-og-salg

[48]Keep in mind that the average applied-for wage does not represent the reservation wage in the baseline model. Instead the reservation wage could be proxied by the lowest applied-for wage in a given time period. In Section A.2 we explicitly take a stance on how applications relate to search behavior in a random search model with reservation wage decisions. In short, we reinterpret the standard job-offer rate as the rate at which a worker successfully learns of a potential job opportunity. Upon learning about the job, the worker now faces a decision about whether to apply for the job or not. This decision is summarized in an (application) reservation wage.

[49]Still, a change in the reservation wage would also affect the relative probability of observing applications to all other jobs above the truncation point. Thus, declining reservation wages may indeed also generate declines in the upper part of the applied-for wage distribution, but only because the set of acceptable job offers has increased changing the relative frequency of specific wage offers in the sample.

Figure 6: Different measures of firm type

(a)

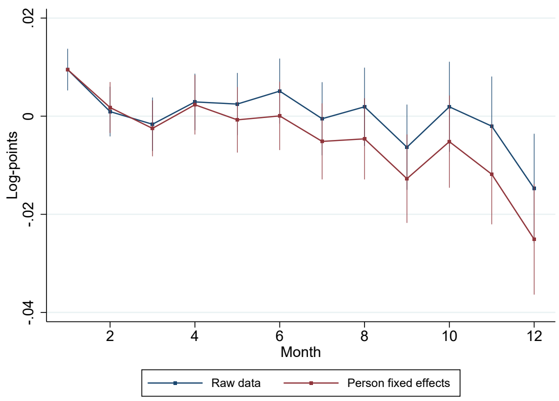

(b) Firm average wage rank

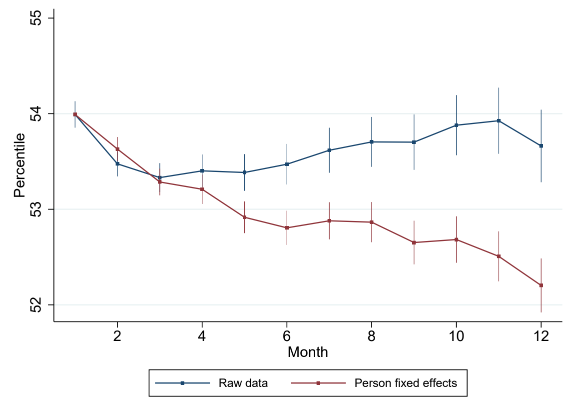

(c) Firm value-added rank

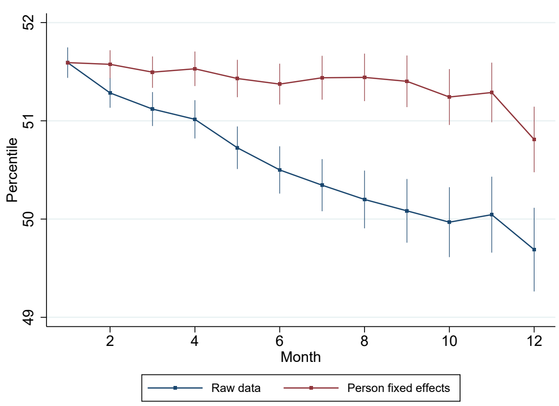

Note: This figure shows the estimates of the duration-month fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, estimated from Equation 6 (blue line) and Equation 7 (red line). Note that in the red line we have normalized τ_1 to equal the mean of the outcome variable in the first month of unemployment. The outcome variables are different measures proxing the firm's "type" as measured by the AKM firm fixed effect, average wages across all employment spells in the firm or the firms value added per worker (see Section C.1). Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

In contrast, in models where job seekers can target their search (for example by choosing between searching in different markets) and thus directly affect which types of job offers they receive, we could see changes throughout the distribution of applications due to changes in the target. Because our data contain several applications per month for each job seeker, it is possible for us to shed some light on this distinction by examining changes in the full distribution of applied-for wages over time.

In Figure 7, we examine changes in the distribution of applied-for wages. Using data on the typical wages of all applied-for jobs, for each job seeker we construct measures of the 80th and 20th percentile of applied-for wages in each month and use these percentiles as the outcome variable in regressions like Equation 7.[50] As noted, meaningfully measuring these percentiles is possible, because our data contain several applications for each worker in each month (see also Table 8). Because we tend to observe markedly fewer applications in the very first month of an unemployment spell (see Figure 4), however, we exclude this month from the analysis when examining wage percentiles. We instead use month 2 as the baseline month in this analysis.[51]

Figure 7 reveals that the decrease in average applied-for wages over time is driven by changes throughout the distribution of applied-for jobs; both the 80th and 20th percentile of applied-for wages shows a decrease that is quite similar to the one seen for the average wage. That is, in the data we find that e.g. the decline in the 80th and 20th percentile of applied-for wages display similar changes (in fact the 80th percentile declines slightly more), indicating a shift throughout the whole set of applied-for jobs. Overall we conclude that the gradual decline in average applied-for wage arises through a level shift throughout the applied-for wage distribution.

We interpret this evidence in favor of at least some degree of targeting in search.[52] An immediate way to reconcile the observed pattern would be to allow for some type of targeting of job search as in e.g. Nekoei and Weber (2017) or in models where search strategies are not completely exogenous, but where job seekers choose e.g. what markets to search in (see footnote 18 for some examples).

C.3 Dynamics in scope of search

In Figure 1c in the main text we showed how the scope of search index evolves over time in unemployment. In Figure 8 we show dynamics separately for each of the sub-components in this index, i.e. commuting time, relatedness of applied-for occupation to the previous job and relatedness of

[50] Results are similar if we focus on other percentiles such as the the 90th and the 10th or simply the highest and lowest applied-for wage.

[51] Because workers take some time to learn about and get instruction in the use of the Joblog registration system, we see a clear pattern that the number of registered jobs in the first month is relatively low. On average workers in our analysis register 4.3 jobs in the first month of their unemployment spell, but 5.8 applications per month averaged over the full spell (see Table 8).

[52] We note that this argument we make is similar in spirit to e.g. Garibaldi et al. (2016); Lentz et al. (2021). Here it is argued that a key difference between random and directed/target search models is that in the former the job seeker cannot affect the distribution of job offers he receive above the reservation wage. On the contrary with targeted job search the job seeker can "target" particular parts of the distribution of wages thus leading to differences in e.g. job reallocation patterns. A great advantage of the job search data is that we observe actual job search instead of only the eventual job outcome, this generates ample opportunity to more directly study how targeting plays out.

Figure 7: Changes in the distribution of applied-for wages

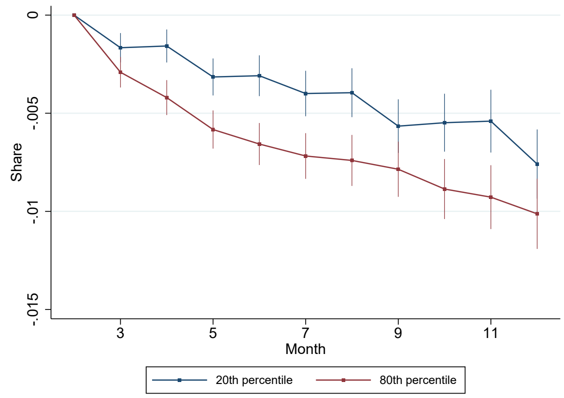

Note: This figure shows the resulting estimates of the duration-month fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, in Equation 7 when the outcome is the 80^{th} (20^{th}) percentile in the distribution of applied-for wages in a given month. Note that in both the blue and the red line we have normalized τ_1 to equal to 0 in the first month of unemployment. Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

applied-for industry. As is clear from the figure, all three components evolve similarly. In particular individuals "broaden" their search along all dimensions, i.e., they are more likely to consider jobs with a longer commute, and jobs which are unrelated to their previous jobs in terms of occupations or industries.

C.3.1 Unrelated occupations

To test the robustness of our results, and in particular our results for the relatedness of occupations, we construct alternative measures of related occupations using the Danish register data. First, we build an occupational-code to occupational-code transition matrix using the Danish BFL register (see Section B.2). We use occupational code changes both within (e.g. promotion) and between jobs (also allowing for shorter periods of unemployment). Explicitly, we follow individuals over time from 2009-2016 and sample all events where the occupational code changes for an employee, allowing for a maximum of 365 days in between payments. Using these transitions we construct a matrix which describes the transition probabilities over future occupational codes conditional on a specific previous occupation code. To be consistent with the definition in the main text (see Section 3.3), we then categorize applied-for jobs in terms of how unrelated they are by focusing on the inverse event (i.e., not making a transition to this specific occupation code).[53] We construct measures at both the 1st, 2nd and 3rd digit levels. Compared to the O*Net based measure, our "BFL measure" has the advantage that we can classify all occupation-to-occupation transitions with their corresponding transition probability, thereby avoiding making choices of how many occupations to classify as (un)related, and further we obtain a richer characterization of how (un)related occupations are.

[53]Note that opposite to the categorization in the main text, the categorization of each applied-for job is no longer binary (unrelated or not) but now instead a probability.

Figure 8: Underlying changes in scope of search

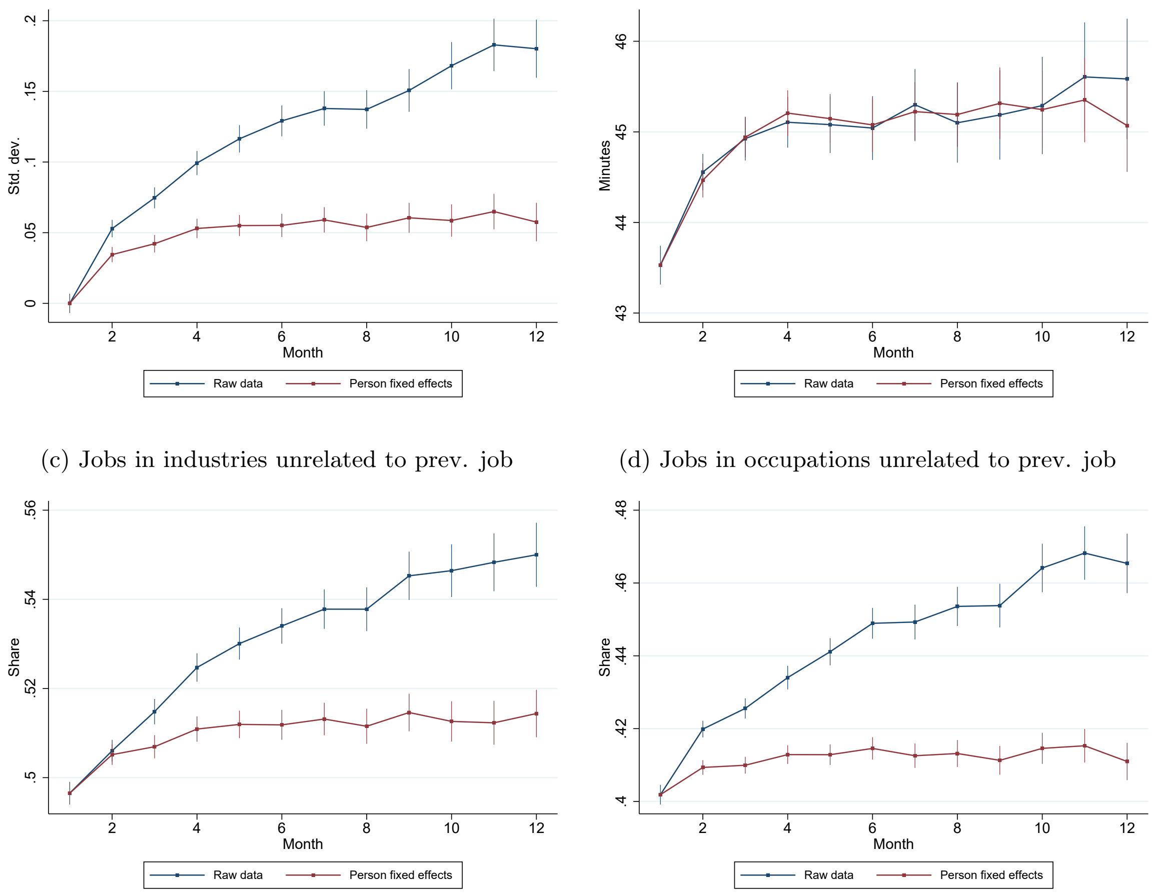

(a) Scope of search

(b) Avg. commute time

(c) Jobs in industries unrelated to prev. job

(d) Jobs in occupations unrelated to prev. job

Note: This figure shows the estimates of the duration-month fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, estimated from Equation 6 (blue line) and Equation 7 (red line). Note that in the red line we have normalized τ_1 to equal the mean of the outcome variable in the first month of unemployment. The outcome variables are the scope of search index and each sub-component for this index. See Section 3.3 for further details on these outcomes. Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

In Figure 9 we report the dynamics in unrelated occupations based on the O*Net matrix (see section B.5) and dynamics in the probability of applying to occupations which are unrelated based on these measure(s). The dynamics are remarkably similar to the dynamics reported in Figure 8, i.e., across all three measures we see a larger initial adjustment during the first 3-4 months, and thereafter the change in estimates decreases substantially, and we see small changes in search behavior in these dimensions.[54] We discuss this further in the next sub-section.

C.3.2 Stock-flow matching and vacancy supply effects

As noted in the main text, changes in scope of search are concentrated entirely in the first three to four months of the unemployment spell. One possible explanation for why job seekers adjust the scope of search of applied-for jobs at the beginning of the unemployment spell is vacancy supply effects, as in the stock-flow model of job search.[55] The idea is that unemployed workers face a large stock of potential vacancies to apply to, but this set quickly narrows as the unemployed worker exhausts the relevant options. This can exactly generate a pattern where workers are pushed to broaden their search over the first few months, but not later on.[56] If the dynamics, and in particular the absence of dynamics after the initial adjustment, in the proximity of applied-for jobs are driven by such vacancy supply effects, we would expect large changes in the probability of applying for vacancies in the stock versus the flow around months 3 to 5. Further it is reasonable to expect applied-for job characteristics to differ based on the duration of the vacancy if we expect vacancies in the stock to partly consist of vacancies which were previously found non-suitable (see also footnote 55).

To provide some evidence on this mechanism, we need a measure of whether each applied-for job is a newly posted vacancy or a vacancy that has been posted for some time. Our data do not contain information on vacancy posting dates. For each applied-for job in our data, however, we can construct a rough proxy of whether the job is likely to be a newly posted vacancy by checking whether the data contain any other applications to the same firm over the preceding two weeks. If no other applications have gone to this firm over the preceding two weeks, we label the job

[54] Note also that the baseline (the average of the outcome variable in month 1) is increasing as we move to higher level occupation codes. This makes sense as for higher digits, the probability of a specific transition decreases, and hence the probability of the inverse is increasing.

[55] In the stock-flow framework, unemployed workers and vacant jobs coexist because they cannot form sufficiently productive matches (i.e. not due to search frictions as such). In its standard form, job search dynamics would arise because the unemployed focus on the inflow of new vacancies since the existing stock of vacancies have already been exhausted. Non-stationarities in the economic environment such as e.g. time-limited benefits or human capital depreciation, may lead the "longer-term unemployed" to partly revisit vacancies in the stock later on in the unemployment spell if preferences over matches change. We note that if this mechanism is present, we should expect applied-for job characteristics to differ depending on whether vacancies are from the flow or the stock.

[56] In existing theoretical models, unemployed workers are assumed to exhaust all existing vacancies in the vacancy stock very early on in the UI spell. This implies that changes in job search behavior due to stock-flow effects should materialize early on in the spell. In practice, it may take some time before a newly unemployed worker has applied to all relevant jobs in the existing vacancy stock. This would imply a more gradual change in behavior over the few months of unemployment. Note that the most common predictions from the stock flow framework would concern the number of submitted applications, which should experience a (large) decline as the unemployed exhaust the stock. Below we take a broader view and consider stock flow effects as driving dynamics in applied-for job characteristics.

Figure 9: Jobs in occupations unrelated to prev. job - alternative measures

(a) Jobs in occupations unrelated to prev. job based on O*Net matrix

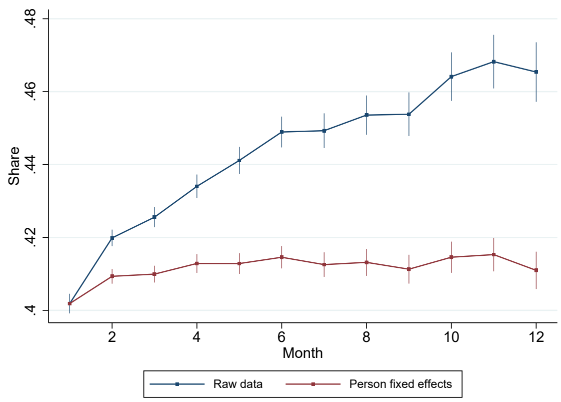

(b) Unrelated occupations based on transition matrix at the 1 digit level

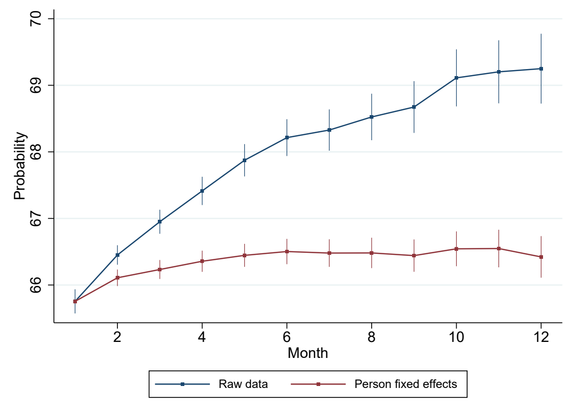

(c) Unrelated occupations based on transition matrix at the 2 digit level

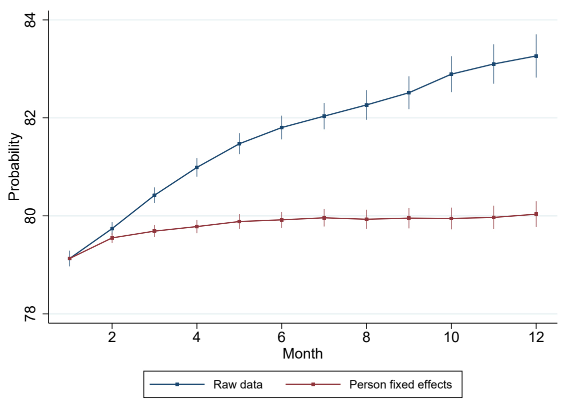

(d) Unrelated occupations based on transition matrix at the 3 digit level

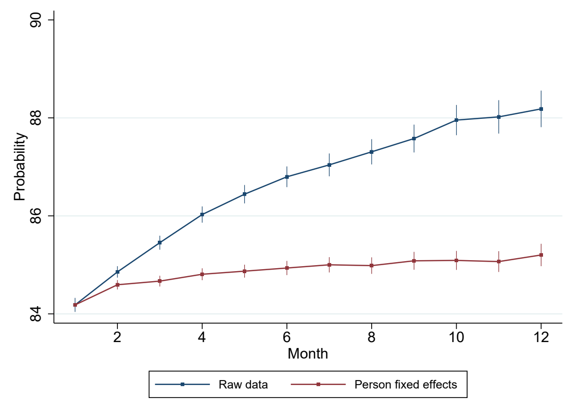

Note: This figure shows the estimates of the duration-month fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, estimated from Equation 6 (blue line) and Equation 7 (red line). The outcome variables are unrelated occupations based on O*Net and alternative measures of unrelated occupations which are created using the Danish registers, see Section C.3.1. Note that we have normalized τ_1 to equal the mean of the outcome variable in the first month of unemployment. Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

as representing a new vacancy and record the vacancy opening week, otherwise we label it as an existing vacancy.

Using this measure of new vacancies, Figure 10 examines how the share of applications going to new vacancies changes over the unemployment spell based on our event study specification.[57] The overall direction of change is consistent with the stock-flow framework: Over time, the share of applications going to new vacancies tends to increase systematically. The size of this increase is small, however; at most one percentage point over the entire period (corresponding to around 3 percent of a standard deviation of the variation across individuals in the first month). More importantly, the timing of the increase is inconsistent with the idea that stock-flow effects are the main explanation for changes in the scope of search of applied-for jobs over the first 4 months of unemployment, since the share of applications going to new jobs is essentially flat for the first 2 months and then slowly increase throughout the remaining spell. Although stock-flow effects may well be at work in the data, they are thus unlikely to explain the observed changes in the proximity of applied-for jobs.

Figure 10: Share of applications going to new vacancies

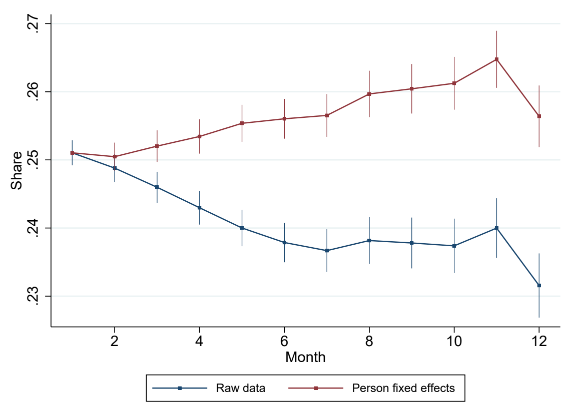

Note: This figure shows the estimates of the duration-month fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, estimated from Equation 6 (blue line) and Equation 7 (red line). Note that in the red line we have normalized τ_1 to equal the mean of the outcome variable in the first month of unemployment. The outcome variable is the share of applications which are submitted to firms that have not received other applications in our data in a 2-week time window preceding the submitted application. Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

C.4 Dynamics in how job was found and how applications are submitted

In Figure 1d in the main text we show how the share of jobs found via vacancy posting evolve over the unemployment spell. This share increases continuously. After accounting for dynamic selection

[57]We define applying to a new vacancy as applying to a vacancy which has opened during the last two weeks (we find similar results when we used other time frames, or if distinguish "vacancies" by the combination of firm id and occupation codes at the 1-digit level.). A brand new vacancy should obviously never have received any applications in the past. Since many firms periodically post new vacancies and hire, however, some time frame needs to be imposed on how far back to look for past applications to the same firm.

(red line), the magnitude of this increase is about 10 percentage points over the first year of the unemployment spell. In this subsection we show that the increased focus on vacancy postings comes both at the expense of applying less to jobs found via social networks and sending less cold contact applications. We also provide some evidence on changes in how the unemployed are applying for jobs throughout the spell.

In Figure 11 we show changes in how applied-for jobs were found throughout the unemployment spell. We find that the increase in jobs found via vacancy posting equally stems from a 4-5 percentage point decrease of applied-for jobs that were found through a personal network or cold contacts. These results may suggest that unemployed workers gradually exhaust informal search methods and substitute for more formal job search, which may also be reflected in the method through which job seekers apply for jobs.

In Figure 12, we report changes in the method used to submit applications throughout the unemployment spell. The figure reveals systematic changes in application channels. Most notably, the share of applications that are submitted through an online form increases markedly throughout the spell. This result has implications for much existing and ongoing research on job search. As discussed in the introduction, much of the existing work using micro data on job search is based on data on applications and search behavior from one or more online job platforms.

The results in Figure 12 show that - at least in the Danish case - job seekers are very systematically selecting into these platforms over time, which raises a number of potential concerns with these types of data. For example, if the number of applications made on an online platform is used to infer the total sent applications, the gradual switch to online platforms shown in Figure 12 will bias the analysis towards finding an increase in applications over time, even if the total number of applications made online and offline is constant over time. For data sets that only include individuals who are active on the search platform, and where the start of a job search spell must be inferred from activity on the platform, the pattern in Figure 12 also suggests that observed samples of "new searchers" can in fact be skewed towards individuals who have been unemployed for an extended period of time.

C.5 Dynamics for posted vacancies only

Figure 1 panel D in the main text shows how the share of applied-for jobs found through posted vacancies increases over the course of unemployment. This could reflect exhaustion of search methods which could affect the dynamics we observe in job characteristics. In Figure 13 we directly compare dynamics for our standard applied-for job characteristics across two samples of applied-for jobs; the full sample of all applied-for jobs and only for jobs "found via publicly posted vacancies". The estimates are based on our specification with individual fixed effects (Equation 7). To facilitate comparisons between the two groups, however, we have also added a version where we normalize the estimated coefficient for month one to be zero for both groups.

In general we find that the dynamics are very similar suggesting that changes in search methods are not key drivers of dynamics in applied-for job characteristics. We do however see that applied-

Figure 11: Changes in how applied-for jobs were found

(a) Jobs found via a publicly posted vacancy

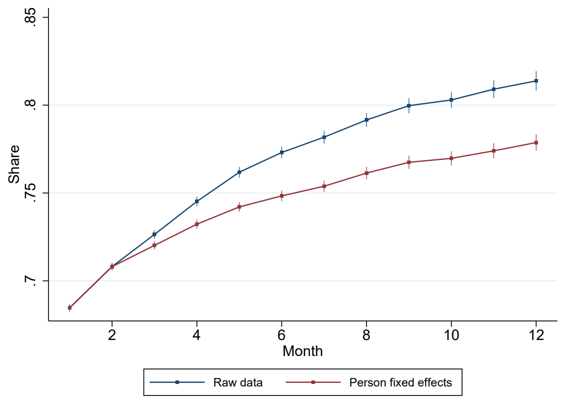

(b) Jobs found via personal network

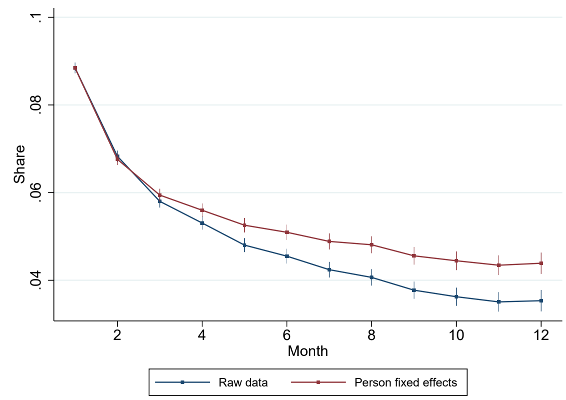

(c) Cold contact applications

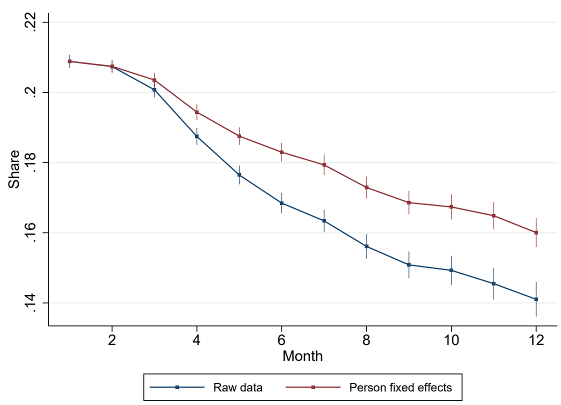

Note: This figure shows the estimates of the duration-month fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, estimated from Equation 6 (blue line) and Equation 7 (red line). Note that in the red line we have normalized τ_1 to equal the mean of the outcome variable in the first month of unemployment. The outcome variable is the share of submitted applications with a specific search method (answers to the question: "how did you find the job?"). Cold contact applications cover applications submitted without knowing whether a vacancy existed, see Section 3.3 for additional details on the different measures. Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

Figure 12: Changes in application channels used to apply

(a) Applications via mail or e-mail

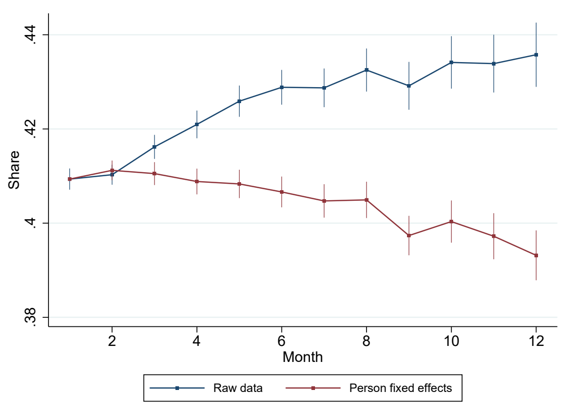

(b) Applications made verbally (phone/in-person)

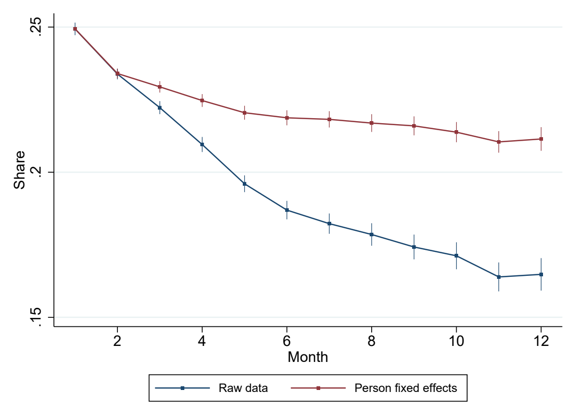

(c) Applications made through an online platform

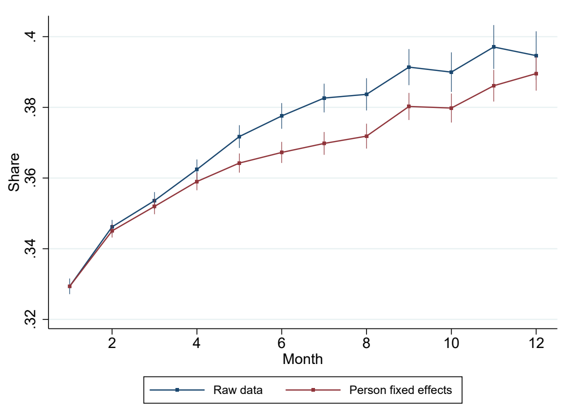

(d) Applications made through other means

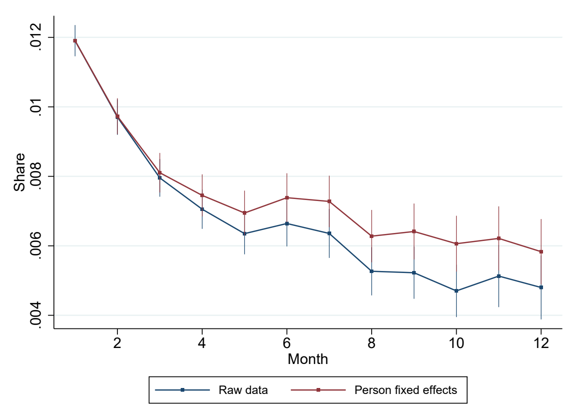

Note: This figure shows the estimates of the duration-month fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, estimated from Equation 6 (blue line) and Equation 7 (red line). Note that in the red line we have normalized τ_1 to equal the mean of the outcome variable in the first month of unemployment. The outcome variable is the share of submitted applications with a specific application method (answers to the question: "how did you apply for the job?"). See Section 3.3 for additional details on the different measures. Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

for jobs "found via publicly posted vacancies" adjust slightly less in our measure of broadness of search compared to the full sample of all applied-for jobs. Evaluating each of the sub-components of the scope of search index (commuting, related occupations and related industries, see Section 3.3) we find that this difference arise because the adjustment path (and initial levels) for related industries differs for posted vs. all applications. We therefore believe a plausible explanation for this discrepancy is industry differences in the use of vacancy posting.

C.6 The role of observables

In Section 4.4 we found that standard observables are capable of explaining a part of the variation in applied-for wages, both overall and especially across medium vs. longer-term unemployed. We now look closer at differences in search behavior for our other outcomes (full-time, posted vacancies and scope of search) across groups of individuals with different realized unemployment duration.

In Tables 13, 14 and 15 we regress measures of search behavior on separate dummies for shorter- and longer-term unemployment spells. The reference group is therefore medium-term unemployed. Across the different columns we change the set of control variables to assess the importance of these variables for the initial level differences we observe across unemployment duration groups. Column (1) in Tables 13, 14 and 15 reports the (raw) differences in search behavior across the three groups when no other observables are included. We see that - compared to medium-term unemployed - a larger share of long-term unemployed's applied-for jobs are for full-time jobs and to jobs that are posted. At the same time, long-term unemployed also have a broader scope of search than medium-term unemployed. Individuals who are unemployed for a shorter-term apply less to jobs that are full-time and less to jobs that are posted, and they also have a narrower scope of search compared to medium-term unemployed.

Broadly speaking these differences appear unrelated to standard observables as none of the added controls close the gap between short- and medium-term unemployed nor between medium- and long-term unemployed. For example, the initial gap for short term unemployed of more than -0.7 of a standard deviation in the scope of search index remains even when the full set of observables is added to the regression. In some other specifications, observables matter more: For example, adding the full set of observables halves the gap in applied-for jobs to posted vacancies for medium- and long-term unemployed (see Table 15). Still, across all outcomes and specifications, a very substantial part of differences in job search behavior across unemployment groups appears to be unrelated to standard observables. It is also worth highlighting that along these dimensions of search behavior, standard observables play a limited role in the overall variation in the first month of search as the R^2 ranges between 5-9 percentage points.

Figure 13: Dynamics for posted vacancies

(a) Average applied-for wage

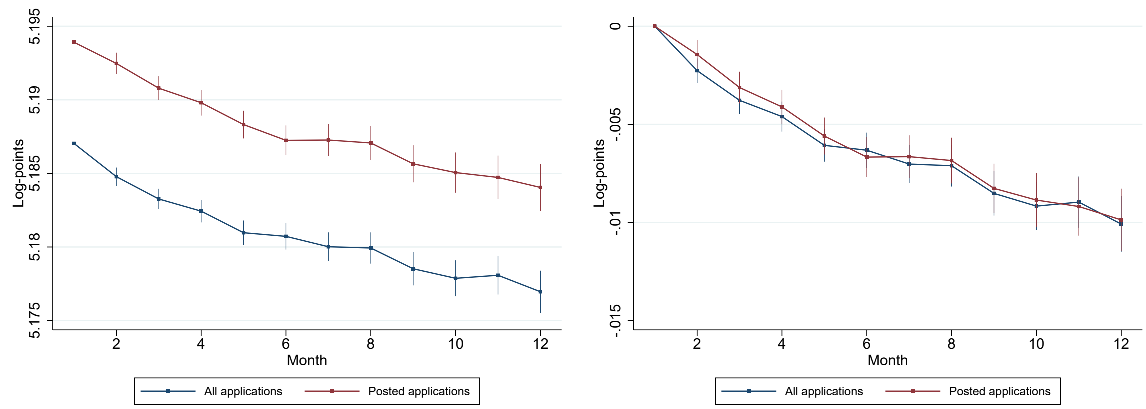

(b) Share of jobs that are full-time

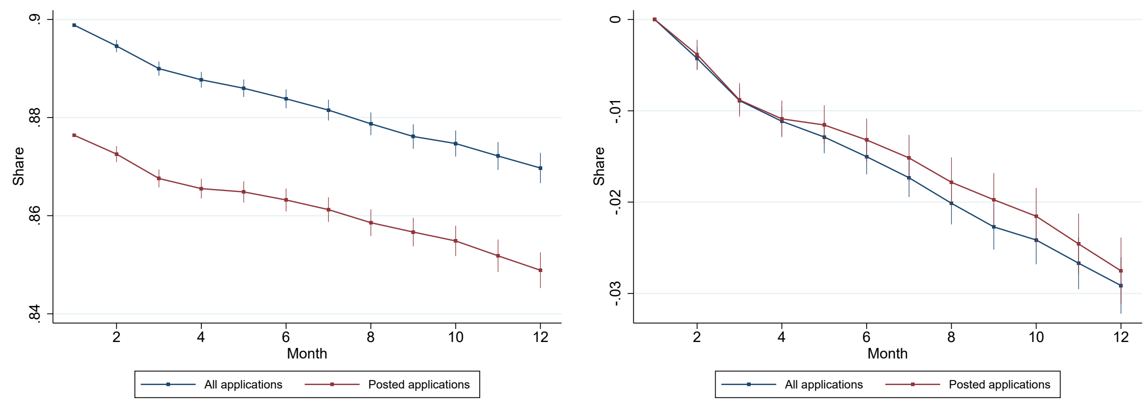

(c) Scope of search

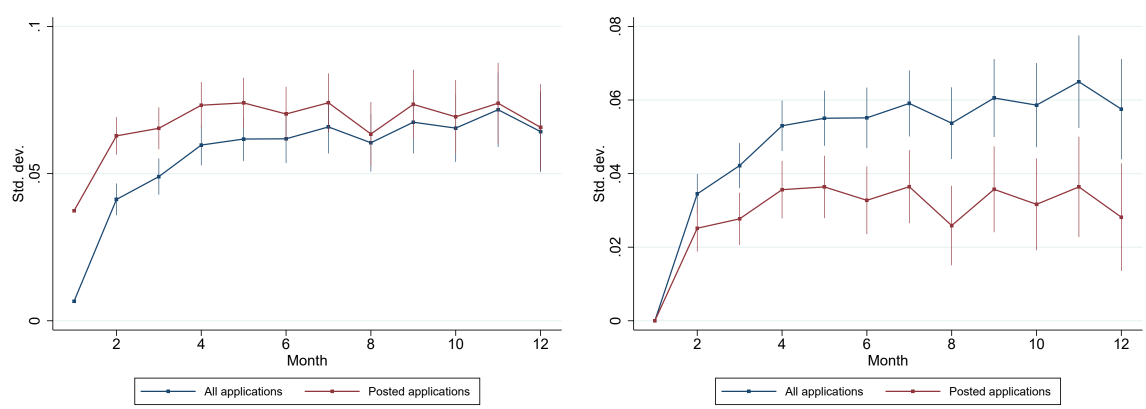

Note: This figure shows the estimates of the duration-month fixed effects, $\tau_1, \tau_2, ...\tau_{12}$, from Equation 7 separately for the full sample of applied-for jobs (blue line) and only for jobs found via posted vacancies (red line). Note that in the left panel we have normalized τ_1 to equal the mean of the outcome variable in the first month of unemployment in each sample. In the right panel the blue and the red lines we have normalized τ_1 to 0 in the first month of unemployment. The outcome variables are the average typical wage of applied-for jobs, share of applications to full-time jobs and the scope of search. See Section 3.3 for further details on these outcomes. Standard errors are clustered at the level of the individual, and vertical bars display 95% confidence bands.

Table 13: Share of jobs that are full-time

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Duration: 2-3 months	-0.00143	-0.00338**	-0.00496***	-0.000463	-0.00102	-0.00108	-0.00171	-0.00629***
	(0.00161)	*(0.00157)*	*(0.00162)*	*(0.00161)*	*(0.00161)*	*(0.00161)*	*(0.00161)*	*(0.00158)*
Duration: >10 months	0.00870***	0.00781***	0.00958***	0.00882***	0.00899***	0.00832***	0.0101***	0.00934***
	(0.00235)	(0.00229)	*(0.00234)*	*(0.00234)*	*(0.00234)*	*(0.00235)*	*(0.00235)*	(0.00228)
Female		-0.0961***						-0.0940***
		(0.00164)						*(0.00165)*
Age		0.000552***						2.36e-05
		(5.89e-05)						*(6.73e-05)*
One or more children		-0.0159***						-0.0216***
		(0.00201)						(0.00202)
Log UI payout			0.0874***					0.0652***
			(0.00450)					*(0.00468)*
Log UI payout/previous earnings				-0.0145***				
				(0.000782)				
Log previous earnings					0.0158***			0.0108***
					(0.000749)			*(0.000845)*
Upper secondary						-0.0176***		-0.00652***
						(0.00231)		*(0.00227)*
Short tertiary/Bachelor/Master/Doctoral						-0.00732***		0.0188***
						(0.00233)		(0.00233)
Any UI experience past year							0.0181***	0.0113***
							(0.00338)	*(0.00330)*
Any UI experience past two years							0.0148***	0.00799***
							(0.00228)	*(0.00222)*
Constant	0.900***	0.934***	0.179***	0.895***	0.765***	0.910***	0.895***	0.317***
	(0.00116)	*(0.00270)*	*(0.0371)*	*(0.00118)*	*(0.00652)*	*(0.00219)*	*(0.00121)*	*(0.0371)*
Observations	87,076	87,076	87,076	87,076	87,076	87,076	87,076	87,076
R-squared	0.000	0.055	0.005	0.004	0.005	0.001	0.002	0.063

Note: This table reports estimates from regressions where the dependent variable is the share of applied-for jobs that are full-time in the first month on unemployment. In all regressions we include dummy variables distinguishing spells by their "completed" unemployment duration. We distinguish between short, medium and long unemployment spells (medium spells are the omitted category). Short spells are 3 months or less, long spells are more than 10 months long. Censored spells with a duration of less than 11 months are excluded from the analysis. Across the different columns we change the set of control variables to assess the importance of these variables in explaining the initial level differences we observe across unemployment duration groups. For further information on control variables, see Appendix B.2.

Table 14: Scope of search

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Duration: 2-3 months	-0.0739***	-0.0876***	-0.0569***	-0.0853***	-0.0789***	-0.0724***	-0.0746***	-0.0786***
	(0.00836)	*(0.00830)*	*(0.00839)*	*(0.00827)*	*(0.00826)*	*(0.00836)*	*(0.00836)*	*(0.00827)*
Duration: >10 months	0.0826***	0.0934***	0.0800***	0.0797***	0.0788***	0.0795***	0.0817***	0.0737***
	(0.0124)	(0.0123)	*(0.0123)*	*(0.0122)*	*(0.0122)*	*(0.0124)*	*(0.0124)*	(0.0122)
Female		-0.160***						-0.203***
		(0.00866)						*(0.00868)*
Age		-0.00578***						0.000567
		(0.000314)						*(0.000360)*
One or more children		-0.161***						-0.100***
		(0.0104)						(0.0104)
Log UI payout			-0.442***					-0.110***
			(0.0236)					*(0.0250)*
Log UI payout/previous earnings				0.165***				
				(0.00416)				
Log previous earnings					-0.163***			-0.158***
					(0.00397)			*(0.00468)*
Upper secondary						-0.0851***		-0.0581***
						(0.0122)		*(0.0121)*
Short tertiary/Bachelor/Master/Doctoral						0.0147		0.0104
						(0.0123)		(0.0124)
Any UI experience past year							0.0742***	0.0380**
							(0.0176)	*(0.0174)*
Any UI experience past two years							-0.0864***	-0.0481***
							(0.0116)	*(0.0114)*
Constant	0.0161***	0.370***	3.664***	0.0722***	1.414***	0.0468***	0.0281***	2.418***
	(0.00601)	*(0.0146)*	*(0.194)*	*(0.00610)*	*(0.0346)*	*(0.0116)*	*(0.00633)*	*(0.197)*
Observations	65,890	65,890	65,890	65,890	65,890	65,890	65,890	65,890
R-squared	0.003	0.023	0.008	0.026	0.028	0.005	0.004	0.045

Note: This table reports estimates from regressions where the dependent variable is the average scope of search index of applied for jobs in the first month on unemployment. In all regressions we include dummy variables distinguishing spells by their “completed” unemployment duration. We distinguish between short, medium and long unemployment spells (medium spells are the omitted category). Short spells are 3 months or less, long spells are more than 10 months long. Censored spells with a duration of less than 11 months are excluded from the analysis. Across the different columns we change the set of control variables to assess the importance of these variables in explaining the initial level differences we observe across unemployment duration groups. For further information on control variables, see Appendix B.2.

Table 15: Share of applications to jobs found via posted vacancy

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Duration: 2-3 months	-0.0261***	-0.0247***	-0.0251***	-0.0270***	-0.0265***	-0.0264***	-0.0240***	-0.0222***
	(0.00272)	*(0.00269)*	*(0.00274)*	*(0.00272)*	*(0.00272)*	*(0.00265)*	*(0.00269)*	*(0.00264)*
Duration: >10 months	0.0291***	0.0309***	0.0289***	0.0290***	0.0289***	0.0235***	0.0203***	0.0165***
	(0.00396)	(0.00392)	*(0.00396)*	*(0.00396)*	*(0.00396)*	*(0.00386)*	*(0.00392)*	(0.00381)
Female		0.116***						0.0936***
		(0.00280)						*(0.00276)*
Age		-0.00125***						0.000366***
		(0.000101)						*(0.000112)*
One or more children		-0.000778						-0.00270
		(0.00344)						(0.00338)
Log UI payout			-0.0243***					0.0142*
			(0.00762)					*(0.00782)*
Log UI payout/previous earnings				0.0130***				
				(0.00132)				
Log previous earnings					-0.0126***			0.00204
					(0.00127)			*(0.00141)*
Upper secondary						0.0853***		0.0692***
						(0.00380)		*(0.00379)*
Short tertiary/Bacehlor/Master/Doctoral						0.229***		0.196***
						(0.00383)		(0.00390)
Any UI experience past year							-0.160***	-0.122***
							(0.00564)	*(0.00551)*
Any UI experience past two years							-0.0604***	-0.0464***
							(0.00380)	*(0.00371)*
Constant	0.679***	0.665***	0.880***	0.684***	0.788***	0.546***	0.704***	0.387***
	(0.00195)	*(0.00463)*	*(0.0628)*	*(0.00200)*	*(0.0110)*	*(0.00360)*	*(0.00203)*	*(0.0620)*
Observations	87,076	87,076	87,076	87,076	87,076	87,076	87,076	87,076
R-squared	0.003	0.029	0.003	0.004	0.004	0.055	0.028	0.085

Note: This table reports estimates from regressions where the dependent variable is the share of applications to jobs found via posted vacancy in the first month on unemployment. In all regressions we include dummy variables distinguishing spells by their "completed" unemployment duration. We distinguish between short, medium and long unemployment spells (medium spells are the omitted category). Short spells are 3 months or less, long spells are more than 10 months long. Censored spells with a duration of less than 11 months are excluded from the analysis. Across the different columns we change the set of control variables to assess the importance of these variables in explaining the initial level differences we observe across unemployment duration groups. For further information on control variables, see Appendix B.2.

D Additional discussion of theory

D.1 Model heterogeneity and the two-way fixed effects model

In this section we provide one example of how to modify our benchmark search model so it is consistent with the finding that applied-for wages follow a two-way fixed effects model with individual and time fixed effects. This follows the common approach of introducing scalar heterogeneity across workers with a proportional value of unemployment.

We start from exactly the same model setup as in Section 2, however, we now explicitly introduce heterogeneity across both workers and over time by adding subscripts i and t to model parameters, value functions and choice variables. Specifically, at t months into her unemployment spell, we assume that worker i faces a flow utility of unemployment of b_{it}, a job offer arrival rate of λ_{it}, has a continuation value of unemployment of U_{it} and searches according to a reservation wage of w_{it}^R and an applied-for wage w_{it}^*. The Bellman equation is as before with the exception that changes in the value of unemployment over time, $\frac{dU_{it}}{dt}$, now enters the Bellman equation:

$$rU_{it} = b_{it} + \lambda_{it} \int_0^\infty \max\left(\frac{w}{r} - U_{it}, 0\right) dF(w) + \tfrac{dU_{it}}{dt} \tag{14}$$

We now invoke a scalar heterogeneity assumption and assume that workers differ only in terms of an ability parameter, $a_i \geq 1$, that in turn affects their flow value of unemployment and their offer arrival rate. Letting b_t and λ_t denote the flow utility of unemployment and offer arrival rate at time t for the lowest ability individual (with $a_i = 1$), we specifically assume:

$$b_{it} = b_t \cdot a_i^\rho \tag{15}$$

$$\lambda_{it} = \lambda_t \cdot a_i^\pi \tag{16}$$

In other words we assume that high ability individuals have higher flow utility of unemployment - often interpreted as reflecting a higher value of home production - and that they are more successful at attracting offers. The exogenous parameters, $\rho \geq 0$ and $\pi \geq 0$ govern how rapidly the flow utility and offer rate increase with ability.

To show the link to Equation 9 we need to assume a wage distribution. We assume that F is Pareto with shape parameter $\gamma > 1$ and scale parameter $k > 0$. The Bellman equation can then be simplified to:

$$rU_{it} = b_{it} + \lambda_{it} \left(\frac{k}{rU_{it}}\right)^\gamma \left(\frac{1}{\gamma - 1}\right) U_{it} + \tfrac{dU_{it}}{dt} \tag{17}$$

Now let $\tilde{U_t}$ denote the value function for the lowest ability individual (the solution to Equations 15, 16 and 17when $a_i = 1$) and add the parameter restriction $\rho = 1, \pi = \gamma$. It is easily verified that the following is a solution to the Bellman equation for an arbitrary worker i with ability a_i:

$$U_{it} = a_i \tilde{U}_t$$

As we discuss in Section 5.4, this proportionality in turn implies that the average applied-for wage follows a two-way fixed effects model.