# Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice

MICHÈLE BELOT

*European University Institute and University of Edinburgh*

PHILIPP KIRCHER

*European University Institute and University of Edinburgh*

PAUL MULLER

*University of Gothenburg*

We develop and evaluate experimentally a novel tool that redesigns the job search process by providing tailored advice at low cost. We invited jobseekers to our computer facilities for twelve consecutive weekly sessions to search for real jobs on our web interface. For one-half, instead of relying on their own search criteria, we use readily available labour market data to display relevant alternative occupations and associated jobs. The data indicate that this broadens the set of jobs they consider and increases their job interviews especially for participants who otherwise search narrowly and have been unemployed for a few months.

*Key words*: Online job search, Occupational breadth, Search design

*JEL Codes*: D83, J62, C93

## 1. INTRODUCTION

Getting the unemployed back into work is an important policy agenda and a mandate for most employment agencies. In most countries, one important tool is to impose requirements on benefit recipients to accept jobs beyond their occupation of previous employment, at least after a few months.[1] Yet there is little guidance how they should obtain such jobs and how one might advise them in the process. This reflects the large literature on active labour market policies which is predominantly silent about the effective provision of job search advice, where most studies do not distinguish between advice and enforcement. In their meta-study on active labour market

---

1. See Venn (2012) for an overview of requirements across Organisation for Economic Co-operation and Development countries.

---

*The editor in charge of this paper was Aureo de Paula.*

1

policies, Card *et al.* (2010) merge "job search assistance or sanctions for failing to search" into one category.[2] Ashenfelter *et al.* (2005) assert a common problem that experimental designs "combine both work search verification *and* a system designed to teach workers how to search for jobs" so that it is unclear which element generates the documented success. Only few studies, reviewed in the next section, have focused exclusively on providing advice, mostly through labour-intensive counselling on multiple aspects of job search. Our study aims to contribute by providing and assessing low-cost, automated occupational advice to jobseekers.

Even before evaluating the effects of advice on job search, a prime order question is what advice should be provided and how? In most countries, the provision of advice is usually done by trained advisors who meet jobseekers on a regular basis, yet financial constraints often mean that such advice can only be limited in scope. Our first contribution is to propose an innovative low-cost way of providing tailored advice to jobseekers online. It has long been argued that occupational information is something jobseekers have to learn.[3] Recent evidence both for the U.S. and the U.K. shows a pronounced occupational mismatch (Sahin *et al.*, 2014; Patterson *et al.*, 2016): jobseekers search in occupations with relatively few available jobs while at the same time other occupations with relatively more jobs are available but attract little interest. This "mismatch" has seen a further persistent increase since the great recession. Incomplete information could be a contributor if jobseekers do not fully know which occupations currently have favourable conditions and whether their skills allow them to transit there. The tool we propose aims to address this by suggesting occupations (and shows the jobs that are currently available in them) using an algorithm based on representative labour market statistics. In a nutshell, it recommends additional occupations in which relevant other jobseekers have successfully found jobs and where skill transferability is high, and visualizes where market tightness is favourable.

Our second contribution is to evaluate how the advice provided through our tool affects job search behaviour, that is, to see if and how jobseekers adjust their job search strategies in response to the suggestions they receive and whether this affects job interviews. To do this, we conduct a randomized study in a highly controlled and replicable environment. We recruited jobseekers in Edinburgh from local Job Centres and transformed the experimental labouratory into a job search facility resembling those in "Employability Hubs", which provide computer access to jobseekers throughout the city. Participants were asked to search for jobs via our search platform from computers within our labouratory once a week for a duration of 12 weeks. The main advantage of this "field-in-the-lab" approach is that it allows us to obtain a complete picture of the job search process. Not only do we observe participants' activities on the job search platform, such as the criteria they use to search for jobs and which vacancies they consider; but we also collect information via weekly surveys on which jobs they apply for, whether they get interviews and job offers. Furthermore, we also collect information about other search activities that jobseekers undertake outside the job search platform, which is important if one is worried that effects on any one search channel might simply induce shifts away from other search channels. This allows us to have measures of total search effort and total job interviews that include such effects. These are key advantages of this approach that complement the alternatives reviewed in the next section: Studies that rely on data from large online job search platforms typically do not have information on activities outside the job search platform nor on job search success. They also currently lack a randomized design. Studies that use administrative data usually only have information about final outcomes (*i.e.* job found) but know little about the job search process. Because of the logistics

2. See the clarification in (Card *et al.*, 2009, p. 6).

3. For example, Miller (1984), Neal (1999), Gibbons and Waldman (1999), Gibbons *et al.* (2005), Papageorgiou (2014), and Groes *et al.* (2015) highlight implications of occupational learning and provide evidence of occupational mobility consistent with a time-consuming process of gradual learning about the appropriate occupation.

required for our field-in-the-lab setup, our sample is limited to 300 participants. As a 12-week panel this is a large number for experimental work but limited relative to usual labour market studies, with associated limits in terms of power. Since it is the first study on the use of online advice, we found that the advantages warranted this approach.

Most of the literature in labour economics focuses on evaluating interventions that have been designed by policymakers or field practitioners. We add to this tradition here, not only by evaluating a novel labour market intervention, but also by leading the design of the intervention itself, using insights from labour economics to integrate existing labour market data right into a job search platform. To our knowledge, our study is the first to use the expanding area of online search to provide advice by redesigning the jobs search process on the web, and allows for a detailed analysis of the effects on the job search "inputs" in terms of search and application behaviour and the amount of interviews that participants receive.

Internet-based job search is by now one of the predominant ways of searching for jobs (Kuhn and Mansour, 2014). In the U.K. where our study is based, roughly two-thirds of both jobseekers and employers now use the Internet for search and recruiting (Pollard *et al.*, 2012; ONS, 2013). We set up two search platforms for internet-based job search that access the database of live vacancies of Universal Jobmatch, the official job search platform provided by the U.K. Department of Work and Pensions, which features a vacancy count at over 80% of the official U.K. vacancies. One platform replicates "standard" designs where jobseekers themselves decide which keywords and occupations to search for, similar to interfaces used on Universal Jobmatch and other commercial job search sites. The second "alternative" platform provides targeted occupational advice. It asks participants which (target) occupation they are looking for—which often coincides with the occupation of previous employment. Then a click of a button provides them with two lists containing the most related occupations. The first is based on common occupational transitions that people who have worked in the target occupation make. The second contains occupations for which skill requirements are similar to that in the target occupation. Another click then triggers a consolidated query over all jobs that fall in any of these occupations within their geographic area. Participants can also take a look at maps to see in which occupations the ratio of unemployed workers to available jobs is more favourable-but data availability limits this to aggregated occupational groups. The maps provide direct information on the competition for jobs in an occupation, skill transferability provides information on the occupations in which the jobseeker has realistic chances to fulfil the needs of a job opening, and information on successful transitions combines both because successful transitions require the availability of jobs in the new occupation and the skills to secure those jobs. The benefit of this intervention is that it provides job search advice in a highly controlled manner based on readily available statistical information, entails only advice and no element of coercion (participants were free to continue with the "standard" interface if they wanted to), and constitutes a low-cost intervention.

Job search occurs precisely because people lack relevant information that is costly and time-consuming to acquire. The main benefit of the Internet is the ability to disseminate information at lower cost. In our study we investigate the following hypotheses. First, whether it stimulates jobseekers to consider a wider set of occupations than they would normally do. Secondly, whether this broader search leads to more job interviews. While we expect both to be the case for jobseekers that would normally search over a narrow set of occupations, it might not be the case for those that already explore many occupations. In particular, our set of suggestions might be smaller than what they normally consider, thereby making their search more narrow. How this affects job interviews depends on how they re-target their job search effort. We also expect a stronger impact of the intervention on jobseekers that have been unemployed longer, as they might be more open to new suggestions. While these predictions arise naturally, we provide an illustrative theory model that lays out these considerations in Section 6. We test these predictions relative to

the obvious null hypothesis (that there will be no effect of the intervention). This might be the case if the information that we provide is already known to jobseekers or if the real problem is incentives to search rather than information.

All participants searched with the standard interface for the first 3 weeks, which provides a baseline on how participants search in the absence of our intervention. After 3 weeks, half of the participants were offered to try the alternative interface. We report the overall impact on the treatment group relative to the control group. Overall, we find that our intervention exposes jobseekers to jobs from a broader set of occupations, increasing our measure of breadth by 0.2 standard deviations (which corresponds to the broadening that would occur naturally over 3 months of unemployment). Job applications become broader, and the total number of job interviews increases by 44%. These effects are driven predominantly by jobseekers who initially search narrowly. These narrow searchers experience a 2-fold increase in total job interviews (compared with similarly narrow searchers in the control group). Among those, the effects are mostly driven by those with above-median unemployment duration (more than 80 days), for whom the effects on interviews are even larger. Since we collected information on job interviews obtained through other channels, we can assess possible spillovers. We find positive effects for such other channels, which indicates that our information is helpful beyond the search on our particular platform. In fact, the statistical significant impact on job interviews is driven by significantly larger reported interviews due to search outside the lab. This re-enforcing effect is in contrast to crowding-out found in studies on monitoring and sanctions which led to offsetting reductions in non-monitored activities (Van den Berg and Van der Klaauw, 2006).

We find similar overall patterns across a number of robustness checks in terms of empirical specifications. Significance does vary somewhat with the specification and outcome variable. Most robust are the significant increase in the occupational breadth of search and the increase in interviews for initially narrow jobseekers. We do find heterogeneity in effects. For example, initially broad jobseekers become more narrow in their search and we find no sign of increased interviews. This group also uses the new interface less. The fact that some individuals search broader initially—for example, because they have less specialized skills as in Moscarini (2001)—should not be surprising. They might already search in many occupations beyond their most preferred one, so recommendations from the alternative interface may be less valuable. If they do use it and new information ends up moving their perceived skills further apart, they become less broad.[4] The heterogeneity in adoption and impact between different groups provides one reason why overall effects are weaker and lack robustness. While we do not find any significant negative effects of interviews for any subgroup, some point estimates remain economically sizeable. This warrants further analysis and caution, and it might be promising to target advice to particular subgroups such as those who otherwise search narrowly and experienced somewhat longer unemployment. This is particularly interesting because targeting could be included directly into an online advice tool. Moreover, if the effects are positive, either overall or for a targeted subgroup, the near zero marginal costs of our type of intervention should make it an attractive policy tool.[5]

Yet, any of these conclusions needs to be viewed with caution. Apart from concerns about the power of our study, a true cost–benefit analysis would need further evaluation of effects on job finding probabilities as well as on whether additional jobs are of similar quality (*e.g.* pay similarly and can be retained for similar amounts of time). On that point, our approach shares similarities with the well-known audit studies (*e.g.* Bertrand and Mullainathan, 2004) on discrimination. The

---

4. If they are already nearly indifferent between occupations, information can only increase differences. We lay this out in our own toy model that includes learning.

5. Designing the alternative interface cost £20,000, and once this is programmed, rolling it out more broadly would have no further marginal cost of an existing platform such as Universal Jobmatch.

main outcome variable in these studies is usually the employer call-back rate rather than actual hiring decisions. As we elaborate in Section 5, it is evident that our study was not intended to pick up effects on job finding because of its size compared to the very low baseline rate of job finding. We find indeed no indication of increased job finding—even in point estimates (though also no significant difference between the point estimate for job finding and the large positive impact on job interviews). We acknowledge that this might not only be due to power issues, though. For example, the conversion rates of interviews into jobs in broader occupations could be lower.[6] A larger-scale assessment would be necessary here. Moreover, a broader roll-out in different geographic areas would also be needed to uncover any general equilibrium effects, which could reduce the effects if search by some jobseekers negatively affects others, or could boost the effects if firms react to more efficient search with more job creation. Such general equilibrium effects may be important (as highlighted by Crépon *et al.* (2013) and Gautier *et al.* (2018)). We hope that future work with conventional large-scale search providers will marry the benefits of our approach with their large sample sizes.

The essence of our findings can be captured in a simple learning theory of job search that is presented in the pan-ultimate section. It also exposes why narrow individuals with slightly longer unemployment duration might be particularly helped by our intervention. In essence, after losing their job individuals might initially search narrowly because jobs in their previous occupation appear particularly promising. If the perceived difference with other occupations is large, our endorsement of some alternative occupations does not make up for the gap. After a few months, unsuccessful individuals learn that their chances in their previous occupation are lower than expected, and the perceived difference with other occupations shrinks. Now alternative suggestions can render the endorsed occupations attractive enough to be considered. Our intervention then induces search over a larger set of occupations and increases the number of interviews. One can contrast this with the impact on individuals who already search broadly because they find many occupations roughly equally attractive. They can rationally infer that the occupations that we do not endorse are less suitable, and they stop applying there to conserve search effort. Their breadth declines, but effects on job interviews are theoretically ambiguous because search effort is better targeted, which might be the reason for the insignificant effects on job interviews for this group in our empirical analysis.

The subsequent section reviews related literature. Section 3 outlines how our study is set up. Section 4 sets the stage by providing basic descriptives about the job search process and the subject pool, covering also issues of representativeness, sample balance, and attrition. Section 5 assesses the impact of our intervention within our main empirical specification as well as in a number of robustness checks. Section 6 uses a simple model to illustrate the forces that might underlie our findings, and the final section concludes.

## 2. RELATED LITERATURE

As mentioned, most studies on job search assistance evaluate a combination of advice and monitoring/sanctions. An example in the U.K. is the work by Blundell *et al.* (2004) that evaluates the New Deal for the Young Unemployed. The programme instituted bi-weekly meetings with a personal adviser to "encourage/enforce job search". The authors establish significant impact of the programme, but cannot distinguish whether "assistance or the "stick" of the tougher monitoring played the most important role" (p. 601). More recently, Gallagher *et al.* (2015) of the

---

6. For example, Moscarini (2001) outlines a model where those who search narrow have particular advantages in those narrow sectors which would not extend equally to search over a broader set of occupations. This might be reflected only in lower interview rates, but could conceivably also affect the conversion rates.

U.K. government's Behavioural Insights Team undertook a randomized trial in Job Centres that re-focuses the initial meeting on search planning, introduced goal-setting but also monitoring, and included resilience building through creative writing. They find positive effects of their intervention, but cannot attribute it to the various elements.[7] Their study indicates the potential of additional information provision.

Despite the fact that a lack of information is arguably one of the key frictions in labour markets, we are only aware of a few studies that exclusively focus on the effectiveness of information interventions in the labour market.[8] Prior to our study the focus has been on the provision of counselling services by traditional government agencies and by new entrants from the private sector. Behaghel *et al.* (2014) and Krug and Stephan (2013) provide evidence from France and Germany that public counselling services are effective and outperform private sector counselling services. The latter appear even less promising when general equilibrium effects are taken into account (Crépon *et al.*, 2013). Bennmarker *et al.* (2013) find overall effectiveness of both private and public counselling services in Sweden. The upshot of these studies is their larger scale and the access to administrative data to assess their effects. The downside is the large costs that range from several hundred to a few thousand euros per treated individual, the multi-dimensional nature of the advice, and the resulting "black box" of how it exactly affects job search. Our study can be viewed as complementary. It involves nearly zero marginal cost, the type of advice is clearly focused on occupational information, it is standardized, its Internet-based nature makes it easy to replicate, and the detailed data on actual job search allow us to study the effects not only on outcomes but also on the search process.

Contemporaneously, Altmann *et al.* (2018) analyse the effects of a brochure that they sent to a large number of randomly selected jobseekers in Germany. It contained information on (1) labour market conditions, (2) duration dependence, (3) effects of unemployment on life satisfaction, and (4) importance of social ties. They find some positive impact, but only for those at risk of long-term unemployment. In our intervention we also find the strongest effects for individuals with longer unemployment duration, but even overall effects are significant and occur much closer in time to the actual provision of information. Their study has low costs of provision and is easily replicable. On the other hand, it is not clear which of the varied elements in the brochure drives the results, there are no intermediate measures on how it affects the job search process, and the advice is generic to all jobseekers.

Our study is also complementary to a few recent studies which analyse data from commercial online job boards. Kudlyak *et al.* (2014) analyse U.S. data from Snagajob.com and find, among other things, that jobseekers lower their aspirations over time. Faberman and Kudlyak (2014) analyse the same data source and show that there is little evidence that declining search effort causes the declining job finding rate. The data lack some basic information such as (un)employment status, but they document some patterns related to our study: Occupational job search is highly concentrated and any exogenous intervention absent it broadens significantly but only slowly over time.

---

7. This resembles findings by Launov and Waelde (2013) that attribute the success of German labour market reforms to service restructuring (again both advice and monitoring/sanctions) with non-experimental methods.

8. There are some indirect attempts to distinguish between advice and monitoring/sanction. Ashenfelter *et al.* (2005) cite experimental studies from the U.S. by Meyer (1995) which have been successful but entailed monitoring/sanctions as well as advice, and they then provide evidence from other interventions that monitoring/sanctions are ineffective in isolation. This leads them to conclude indirectly that the effectiveness of the first set of interventions must be due to the advice. Yet subsequent research on the effects of sanctions found conflicting evidence: for example, Micklewright and Nagy (2010) and Van den Berg and Van der Klaauw (2006) also find only limited effects of increased monitoring, while other studies such as Van der Klaauw and Van Ours (2013), Lalive *et al.* (2005) and Svarer (2011) find strong effects.

Marinescu and Rathelot (2018) investigate differences in market tightness as a driver of aggregate unemployment using data from Careerbuilder.com and concur with earlier work that differences in market tightness are not a large source of unemployment. In their data set search is rather concentrated, with 82% of applications staying in the same city.[9] Using the same data source, Marinescu (2017) investigates equilibrium effects of unemployment insurance. Marinescu and Wolthoff (2016) use data from Careerbuilder.com and document that job titles are informative above and beyond wage and occupational information for attracting applications. As mentioned, these studies have large sample size and ample information of how people search on the particular site, but none involves a randomized design nor do they have information on other job search channels. Also, their focus is not on advice.

Our weekly survey of job search activity outside the lab is related to the seminal study by Krueger and Mueller (2016) that conducted weekly interviews regarding reservation wages with a panel of jobseekers in the U.S. over the course of half a year. Our recommendation to target occupational information to jobseekers that otherwise search narrowly is in the spirit of recent discussion of profiling in active labour market policy. Profiling singles out subsets of individuals for treatment according to a probabilistic assessment of the benefits (see, *e.g.*, Berger *et al.* (2000) for a comprehensive discussion). Interestingly, in our environment the profiling could be integrated directly into a standard job search engine in which individuals first search "normally" and subsequently, depending on the breadth of their search, occupational information could be offered or not. To our knowledge, our study is the first that undertakes job search platform design and evaluates it. While the rise in Internet-based search will render such studies more relevant, the only other study of search platform design that we are aware of is Dinerstein *et al.* (2018), who study a change of the presentation of search results at the online consumer platform eBay.

## 3. THE SET-UP OF THE STUDY

Two main contributions underlie our study: first, we design of a novel online tool that provides labour market information which is readily available to researchers but usually not to jobseekers. The aim is to make this available in an easily accessible cost-effective form and to enable a direct link to potential jobs. Secondly, we evaluate the new tool experimentally in a randomized controlled experiment for which we invited jobseekers for a period of 12 weeks. We used a "standard" interface for comparison. We now describe the experimental design and provide descriptives on the sample and the job search process.

### 3.1. *Description of the advice interface*

We designed an online job search interface in collaboration with professional programmers from the IT Applications Team at the University of Edinburgh. The main feature of the interface is to provide a tailored list of suggestions of possible alternative occupations that may be relevant to job-seekers, based on a preferred occupation that jobseekers pre-specify (but can change at any time).

We use two methodologies to compile a list of suggested alternative occupations. The first methodology builds on the idea that successful labour market transitions experienced by people with a similar profile contain useful information about occupations that may be suitable alternatives to the preferred occupation: the fact that others found jobs there indicates that skills might be transferable and jobs available. It is based on the standard idea in the experimentation literature that others have already borne the cost of experimentation and found suitable outcomes.

---

9. Based on Figure 5 in the 2013 working paper. The "distaste" for geographical distance backed out in this work for the U.S. is lower than that backed out by Manning and Petrongolo (2017) from more conventional labour market data in the U.K., suggesting that labour markets in the U.K. are even more local.

To do this, we use information on labour market transitions observed in the British Household Panel Survey and the national statistical database of Denmark (because of larger sample size).[10] Both databases follow workers over time and record in what occupation they are employed. We then match the indicated preferred occupation to the most common occupations to which people employed in the preferred occupation transition to.[11] This methodology has the advantage of being highly flexible and transportable. Many countries now have databases that could be used to match this algorithm. That is, the tool we propose can easily be replicated and implemented in many different settings.

The second methodology uses information on transferable skills across occupations from the U.S. based website O*net, which is an online "career exploration" tool sponsored by the U.S. department of Labour, Employment & Training Administration. For each occupation, they suggest up to ten related occupations that require similar skills. We retrieved the related occupations and presented the ones related to the preferred occupation as specified by the participant. This provides information on skill transferability only, not on job availability.

The tool was directly embedded in the job search interface. That means that once participants had specified their preferred occupation, they could click "Save and Start Searching" and were taken to a new screen where a list of suggested occupations was displayed. The occupations were listed in two columns. The left column suggests occupations based on the first methodology (labour market transitions) and the right column suggests occupations based on the second methodology (O*net-related occupations). A screenshot is presented in Online Appendix 8.5.6. Participants were informed of the process by which these suggestions came about, and could select or unselect the occupations they wanted to include or exclude in their search. By default all were selected. By clicking the "search" button, the program searched through the same underlying vacancy data as in the control group.[12]

In addition, the interface provides visual information on the labour market tightness for broad occupational categories across regions in Scotland. The goal here is to provide information about how competitive the labour market is for a given set of occupations—which is closest to the idea of search mismatch in Sahin *et al.* (2014) and provides information on the competition for jobs but not on skill transferability. We constructed "heat maps" that use labour market statistics for Scotland and indicate visually (with a colored scheme) where jobs may be easier to get because there are many jobs relative to the number of interested jobseekers. These maps were created for each broad occupational category (two-digit Standard Occupational Classification (SOC) codes).[13] Participants could access the heat maps by clicking on the button "heat map" which was available for each of the suggested occupations based on labour market transitions.

In principle this tool can be used with any database of vacancies that includes occupational codes; for our experimental approach we combine it with one of the largest databases in the U.K.

---

10. The name of the database is IDA—Integrated Database for Labour Market Research administered by Statistics Denmark. We are grateful to Fayne Goes for providing us with the information.

11. For each occupation, we created a list of three to five common transitions using information from both data sets. The list contained all occupations that occur in the top-ten common transitions in both data sets (if there were more than five of these, we selected the five highest occurring occupations). In case this resulted in less than three occupations, we added the highest ranked transitions from each data set until the list contained at least three occupations.

12. Occupations in O*net have a different coding and description and have a much finer categorization than the three-digit occupational code available in the British Household Panel Survey (BHPS) and in Universal Jobmatch. We therefore asked participants twice for their preferred occupation, once in O*net form and BHPS form. The query on our vacancy database relies on keyword search, taking the selected occupations as keywords, to avoid problems of differential coding.

13. These heat maps are based on statistics provided by the Office for National Statistics, (NOMIS, claimant count, by occupations and county, see https://www.nomisweb.co.uk/). An example of a heat map is in Online Appendix 8.5.6.

## 3.2. *Control treatment: standard search interface*

We designed a standard job search engine that replicates the search options available at the most popular search engines in the U.K. (such as Monster.com and Universal Jobmatch). As in the treatment group this allowed us to record precise information about how people search for jobs (what criteria they use, how many searches they perform, what vacancies they click on and what vacancies they save), as well as collecting weekly information (via the weekly survey) about outcomes of applications and search activities outside the laboratory.

A screenshot of the standard search interface is provided in Online Appendix 8.5.6. Participants can search using various criteria (keywords, occupations, location, salary, preferred hours), but do not have to specify all of these. Once they have defined their search criteria, they can press the search button at the bottom of the screen and a list of vacancies fitting their criteria will appear. The information appearing on the listing is the posting date, the title of the job, the company name, the salary (if specified) and the location. They can then click on each individual vacancy to reveal more information. Next, they can either choose to "save the job" (if interested in applying) or "not save the job" (if not interested). If they choose not to save the job, they are asked to indicate why they are not interested in the job from a list of possible answers.

As in most job search engines, they can modify their search criteria at any point and launch a new search. Participants had access to their profile and saved vacancies at any point in time outside the laboratory, using their login details. They could also use the search engine outside the laboratory, but this turned out to be only a very small share compared to the search activities performed in the lab.

The key feature of this interface is that jobseekers themselves have to come up with the relevant search criteria, as is shared by commercial and public jobsearch sites at the time of our study. In order for the study to provide a valid environment to study search behaviour, it is important that the platforms are used seriously and are not viewed as inferior to alternative search environments. In an exit survey we asked all users to evaluate the interface and found that it was evaluated very positively. The responses to the question "How would you rate the search interface compared to other interfaces?" were: Poor (7%), Below average (7%), Average (14%), Good (46%), Very Good (26%). These responses were very similar across the two interfaces.

## 3.3. *Vacancies*

In order to provide a realistic job search environment, both search interfaces access a local copy of the database of real job vacancies of the government website Universal Jobmatch. This is a very large job search website in the U.K. in terms of the number of vacancies, which is crucial because results can only be trusted to resemble natural job search if participants use the lab sessions for their actual job search. The large set of available vacancies combined with our carefully designed job search engine assures that the setting was as realistic as possible.

Each week there are between 800 and 1,600 new vacancies posted in our data set in Edinburgh (see Online Appendix 8.1 for further details). Furthermore, there is a strong correlation between vacancy posting in Edinburgh and the U.K. Comparing the number of vacancies in our database with the official national vacancy statistics suggests that our coverage is above 80%, which is a very extensive coverage compared to other online platforms.[14] It is well-known that not all

---

14. For comparison, the largest U.S. jobsearch platform has 35% of the official vacancies; see Marinescu (2017), Marinescu and Wolthoff (2016) and Marinescu and Rathelot (2018). The size difference might be due to the fact that the U.K. platform is run by the U.K. government.

vacancies on online job search platforms are genuine, so the actual number might be somewhat lower.[15] We introduced ourselves a small number of additional posts (below 2% of the database) for a separate research question (addressed in Belot *et al.* (2018)).[16]

### 3.4. *Jobseekers*

To study the effect of information provision through the new interface, we recruited jobseekers in the area of Edinburgh in two waves: wave 1 was conducted in September 2013 and wave 2 in January 2014. Labour market conditions in Edinburgh are broadly consistent with national ones, displaying monotonically decreasing unemployment between 2012 and 2015 for both (see Online Appendix 8.2).

The eligibility criteria for participating to the study were: being unemployed, searching for a job for less than 12 weeks (a criterion that we did not enforce), and being above 18 years old.[17] We imposed no further restrictions in terms of nationality, gender, age or ethnicity. We aimed to recruit 150 participants per wave, which constitutes about 2% of the stock of jobseeker allowance (JSA) claimants.[18]

As a background on the institutional setting, individuals on JSA receive between £52.25 and £72 per week depending on age. Eligibility depends on sufficient contributions during previous employment or on sufficiently low income.[19] This is linked to the requirement to be available and actively looking for work. In practice, this implies committing to agreements made with a work coach at the job centre, such as meeting the coach at regular (usually bi-weekly) intervals, applying to suggested vacancies, or participating in suggested training. They are not entitled to reject job offers because they dislike the occupation or the commute, except that the work coach can grant a period of up to 3 months to focus on offers in the occupation of previous employment, and required commuting times are capped at 1.5 hours per leg. The work coach can impose sanctions on benefit payments in case of non-compliance to any of the criteria.

We obtained the collaboration of several local public unemployment agencies (called Jobcentre Plus) to recruit jobseekers on their premises during a 2-week window prior to each wave. Since most individuals on jobseeker allowance meet their advisers bi-weekly, this gives us a chance to encounter most of them. The counsellors were informed of our study and were asked to advertise the study. We also placed advertisements at public places in Edinburgh (including libraries and

15. For Universal Jobmatch fake vacancies covering 2% of the stock have been reported, posted by a single account (*Channel 4* (2014)) and speculations of higher total numbers of fake jobs circulate (*Computer Business Review* (2014)). Fishing for CV's and scams are common, including on Careerbuilder.com (*The New York Times* (2009)) and Craigslist.

16. Participants were fully informed about this. They were told that "we introduced a number of vacancies (about 2% of the database) for research purposes to learn whether they would find these vacancies attractive and would consider applying to them if they were available". They were asked for consent and were informed if they expressed interest in them before any actual application costs were incurred, so any impact was minimized. This small number is unlikely to affect job search, and there is no indication of differential effects by treatment group: In an exit survey the vast majority of participants (86%) said that this did not affect their search behaviour, and this percentage is not statistically different in the treatment and control group (*p*-value 0.99). This is likely due to the very low numbers of fake vacancies and to the fact that fake advertisements are common (see footnote 15) and discussed in search guidelines (*e.g.*, Joyce, 2015).

17. We drop one participant from the sample because this participant had been unemployed for over 30 years and was therefore an extraordinary outlier. We also exclude two participants who showed up once without searching and never returned. Including them in the analysis has no effects on the qualitative findings.

18. The stock of JSA claimants In Edinburgh during our study is 9,000 with a monthly inflow of 1,800 approximately.

19. Benefits of £56.25 per week apply to those aged up to 24 years, and £72 per week for those aged 25 years and older. Contribution-based JSA lasts for a maximum of 6 months. Afterwards—or in the absence of sufficient contributions—income-based JSA applies, which is identical to weekly benefits but which extra requirements. Once receiving JSA, the recipient is not eligible for income assistance; however, they may receive other benefits such as housing benefits.

cafes) and posted a classified ad on an online platform (Gumtree). Sign up and show up rates are presented in Table 14 in Online Appendix 8.3. Of all participants, 86% were recruited in the Jobcentres. Most of the other participants were recruited through our ad on Gumtree. Out of the visitors at the Jobcentres that we could talk to and who did not indicate ineligibility, 43% signed up. Out of everyone that signed up, 45% showed up in the first week and participated in the study, which is a substantial share for a study with voluntary participation. These figures display no statistically significant difference between the two waves of the study.

We also conducted an online study in which jobseekers were asked to complete a weekly survey about their job search. These participants did not attend any sessions, but simply completed the survey for 12 consecutive weeks. This provides us with descriptive statistics about job search behaviour of jobseekers in Edinburgh and it allows us to compare the online participants with the lab participants. These participants received a £20 clothing voucher for each 4 weeks in which they completed the survey. The online participants were recruited in a similar manner as the lab participants.[20] The sign up rate at the Jobcentres was slightly higher for the online survey (58%). However, only 21% completed the first survey, which is partly caused by one-fourth of the email addresses not being active.

In Section 4.1, we discuss in more detail the representativeness of the sample, by comparing the online and the lab participants with population statistics.

### 3.5. *Experimental procedure*

jobseekers were invited to search for jobs once a week for a period of 12 weeks (or until they found a job) in the computer facilities of the School of Economics at the University of Edinburgh. We conducted sessions at six different time slots, on Mondays or Tuesdays at 10 AM, 1 PM or 3:30 PM. Participants chose a slot at the time of recruitment and were asked to keep the same time slot for the 12 consecutive weeks.

Participants were asked to search for jobs using our job search engine for a minimum of 30 minutes.[21] Afterwards they could continue to search or use the computers for other purposes such as updating their CV or applying for jobs. They could stay up to 2 hours. No additional job search support was offered. Participants received a compensation of £11 per session attended (corresponding to compensation for meal and travel expenses as advized by Jobcentre Plus) and we provided an additional £50 clothing voucher for job market attire for participating in four sessions in a row. [22]

Participants would register in an office at the beginning of each session and were then told to sit at one of the computer desks in the computer laboratory. Before the first session they received a description of the study and a consent form (see Online Appendix 8.5.1). We handed out instructions on how to use the interface (see Online Appendix 8.5.2). We had assistance in the laboratory to answer clarifying questions. Participants were explicitly asked to search as they normally would.

Once they logged in, they were automatically directed to our own website. They were first asked to fill in a survey. The initial survey asked about basic demographics, employment and

---

20. Participants were informed of only one of the two studies, either the on-site study or the online study. The did not self-select into one or the other.

21. The 30-minute minimum was chosen as a trade-off between, on the one hand, minimizing the effect of participation on the natural amount of job search, while on the other hand ensuring that we obtained enough information.

22. Our study did not affect the entitlements or obligations that participants face at the local Jobcentre. All forms of compensation effectively consisted of subsidies, that is, they had no effect on the allowances the jobseekers were entitled to.

TABLE 1
*Randomization scheme*

| | Monday | | | Tuesday | | |
|---|---|---|---|---|---|---|
| | 10 AM | 1 PM | 3:30 PM | 10 AM | 1 PM | 3:30 PM |
| Wave 1 | Control | Treatment | Control | Treatment | Control | Treatment |
| Wave 2 | Treatment | Control | Treatment | Control | Treatment | Control |

unemployment histories as well as beliefs and perceptions about employment prospects, and measured risk and time preferences. From Week 2 onwards, they only had to complete a short weekly survey asking about job search activities and outcomes. For vacancies saved in their search in our facility we asked about the status (applied, interviewed, job offered). We asked similar questions about their search through other channels than our study. The weekly survey also asked participants to indicate the extent to which they had personal, financial, or health concerns (on a scale from 1 to 10). The complete survey questionnaires can be found in the Online Appendices 8.5.4 and 8.5.5.

After completing the survey, the participants were re-directed towards our search engine and could start searching. A timer located on top of the screen indicated how much time they had been searching. Once the 30 minutes were over, they could end the session. They would then see a list of all the vacancies they had saved which could also be printed. This list of printed vacancies could be used as evidence of required job search activity at the Jobcentre. We received no information about the search activities or search outcomes from the Jobcentres. This absence of linkage was important to ensure that jobseekers did not feel that their search activity in our laboratory was monitored by the employment agency. They could then leave the facilities and receive their weekly compensation. Participants could still access our website from home, for example, in order to apply for the jobs they had found.

### 3.6. *Randomization*

All participants used the standard interface in the first 3 weeks of the study. Half of the participants were offered the "alternative" interface, which incorporates our suggestions tool, from Week 4 onwards. Participants were randomized into control (no change in interface) and treatment group (alternative interface) based on their allocated time slot. We randomized the first time slot into treatment and control, and assigned each following time slot in an alternating pattern, to avoid any correlation between treatment status and a particular time slot. Each time slot that was allocated to control (treatment) in the first wave was assigned to treatment (control) in the second wave (see Table 1). The change of interface was not previously announced, apart from a general introductory statement to all participants that included the possibility to alter the search engine over time.

Participants received a written and verbal instruction of the alternative interface (see Online Appendix 8.5.3), including how the recommendations were constructed, in the fourth week. For them, the new interface became the default option when logging on, but it was made clear that using the new interface was not mandatory. Rather, they could simply switch back and forth between interfaces. This ensures that we did not restrict choice, but rather expanded their means of searching for a job.

### 3.7. *Measures of job search*

The main goal of the study is to evaluate how tailored advice affects job search strategies. Our data allow us to examine each step of the job search process: the listing of vacancies to which jobseekers

TABLE 2
*Outcome variables*

| | Listed vacancies | Applications | Interviews |
|---|:---:|:---:|:---:|
| In lab: | | | |
|   Number | ✓ | ✓ | ✓ |
|   Occupational breadth | ✓ | ✓ | |
|   Geographical breadth | ✓ | ✓ | |
|   Core/Non-core occ's | | | ✓ |
| Outside lab: | | | |
|   Number | | ✓ | ✓ |

are exposed, the vacancies they apply to and the interviews they receive. In the weekly survey that participants complete before starting to search, we ask about applications and interviews through channels other than our study. The intervention may affect these outcomes as well, since the information provided in the alternative interface could influence people's job search strategies outside the lab. Of course, ultimately one would also like to evaluate the effects on job finding and the characteristics of the job found (occupation, wage, duration, etc.), which would be important to evaluate the efficiency implications of such an intervention. This is however not the prime goal of this study and given the small sample of participants, we should be cautious when interpreting the results on job finding as we discuss in a separate part in Section 5.5. We summarize the outcome variables of interest in Table 2. The main outcome variables relate to (1) listed vacancies, (2) applications, and (3) interviews. The precise definition of each of these is presented next.

The most immediate measure of search relates to listed vacancies, that is, all vacancies that appear on the participants' screen as a return to their search queries in a given week. For a given search query, up to twenty-five results are presented on the screen and only these are included in the set of listed vacancies.[23] If the query returned more vacancies, the participant can click to move to the next screen where again up to twenty-five vacancies are shown. These are added as "listed" for this week. This means that vacancies are only recorded as "listed" if the applicant has seen them on the screen. All our analyses are at the weekly level and, thus, we group listings from all search queries in a week together.[24] We note that listings are not mechanical even in the treatment group but, rather, remain an outcome of their choice: on the new interface users still decide how many pages of results to move through, which geographical radius to explore, how many recommended alternative occupations to keep, and how many preferred occupations and associated alternatives to explore in a given week—not to mention that participants can revert back to standard keyword search to explore some options more deeply (we document the use of each interface later on).

The second measure of search behaviour relates to applications, which is a more direct measure of interest.[25] For each vacancy that was saved in the laboratory, we asked participants to indicate whether they actually applied to it or not.[26] We can therefore precisely map applications to the timing of the search activity. This is important as there may be a delay between the search and

23. The default ordering of results is by posting date, but alternative orderings can be chosen such as location or salary.

24. Since the alternative interface tends to return more search results (due to the additional suggested occupations), it may necessitate less search queries. For that reason the weekly analysis seems more appropriate compared to results at the level of an individual query. In a given week each vacancy is counted at most once.

25. We also record "viewed vacancies" (vacancies that the jobseeker clicks on in order to view all job details) and "saved vacancies", but we prefer to focus on applications as they constitute a more direct measure of interest. Results for viewed and saved vacancies are reminiscent of those for listed and applied vacancies and are omitted for brevity.

26. If they had not applied, they are asked whether they intended to apply and only if they answered affirmatively they were asked again the next week. A similar procedure was followed for interviews.

the actual application; so applications that are made in week 4 and after could relate to search activity that took place before the actual intervention. For applications conducted based on search outside the laboratory, we do not have such precise information. We asked how many applications jobseekers made in the previous week but we do not know the timing of the search activity these relate to. For consistency, we assume that the lag between applications and search activity is the same inside and outside the laboratory (which is one week) and assign applications to search activity 1 week earlier. As a result, we cannot use information on search activity in the last week of the experiment, as we do not observe applications related to this week.

For listed vacancies and applications we look at the number as well as measures of breadth (occupational and geographical). For occupational breadth we focus on the U.K. SOC code of a particular vacancy, which consists of four digits.[27] The structure of the SOC codes implies that the more digits the two vacancy codes share, the more similar they are. Our measure of diversity within a set of vacancies is based on this principle, defining for each pair within a set the distance in terms of the codes. The distance is zero if the codes are the same, it is 1 if they differ on the last digit, 2 if they differ on the last two digits, etc. This distance, averaged over all possible pairs within a set, is the measure that we use in the empirical analysis (we discuss robustness to alternative measures in Section 5.6). Note that this distance is increasing in breadth (diversity) of a set of vacancies. We compute this measure for the set of listed and applied vacancies in each week for each participant. For geographical breadth we use a simple measure. Since a large share of searches restricts the location to Edinburgh, we use the weekly share of a participant's searches that goes beyond Edinburgh as the measure of geographical breadth.[28]

Our third outcome measure is interviews—which is the measure most closely related to job prospects. As was done for applications, we assign interviews to the week in which the search activity was performed, and assign interviews through channels other than the lab to search activity 2 weeks earlier. As a result we do not use information on search activity in Weeks 11 and 12 of the experiment, because for job search done in these weeks we do not observe interviews. The number of interviews is too small on average to compute informative breadth measures. As an alternative, we asked individuals at the beginning of the study about three "core" occupations in which they are looking for jobs, and we estimate separately the impact of the treatment for interviews in core and non-core occupations.

## 4. DESCRIPTIVE STATISTICS ON JOBSEEKER CHARACTERISTICS AND JOB SEARCH BEHAVIOUR

This section provides descriptive statistics about the characteristics of the sample of jobseekers in our study. We indicate how our experimental sample compares to the (limited) information we have on the overall set of JSA claimants in Edinburgh and to those participating in the online survey, and we demonstrate balance between treatment and control group. The control group faces no intervention throughout the study, and we document how they change their job search over time. For the treatment group we document to which extent they adopt the new interface. Finally, we discuss attrition.

---

27. The first digit defines the "major group", the second the "sub-major group", the third the "minor group" and the fourth the "unit group". Examples are "Social science researchers" (2322) and "Call centre agents/operators" (7211).

28. The surroundings of Edinburgh contain only small towns, the nearest city (Glasgow) takes 1–1.5 hours of commuting.

TABLE 3

*Characteristics of lab participants and online survey participants (based on the first week initial survey)*

|  | Lab participants | | | Online survey | | | T-test[a] | Pop.[b] |
|---|---|---|---|---|---|---|---|---|
|  | Mean | Min | Max | Mean | Min | Max | *p*-value |  |
| Demographics: | | | | | | | | |
| Gender (%) | 43 | | | 52 | | | 0.09 | 33 |
| Age | 36 | 18 | 64 | 34 | 18 | 64 | 0.08 | 35 |
| High educ[c] (%) | 43 | | | 43 | | | 1.00 | |
| White (%) | 80 | | | 77 | | | 0.43 | 89 |
| Number of children | 0.53 | 0 | 5 | 0.28 | 0 | 2 | 0.02 | |
| Couple (%) | 23 | | | 23 | | | 0.96 | |
| Any children (%) | 27 | | | 23 | | | 0.41 | |
| Job search history: | | | | | | | | |
| Vacancies applied for | 64 | 0 | 1000 | 75 | 0 | 1354 | 0.53 | |
| Interviews attended | 2.3 | 0 | 50 | 2.7 | 0 | 20 | 0.42 | |
| Jobs offered | 0.42 | 0 | 8 | 0.51 | 0 | 10 | 0.52 | |
| At least one offer (%) | 20 | | | 24 | | | 0.36 | |
| Days unempl. (mean) | 260 | 1 | 5141 | 167 | 8 | 2929 | 0.15 | 111 |
| Days unempl. (median) | 80 | | | 118 | | | | 81 |
| Less than 183 days (%) | 76 | | | 75 | | | 0.76 | |
| Less than 366 days (%) | 85 | | | 91 | | | 0.13 | |
| Jobseekers allowance (£) | 52 | 0 | 1005 | 58 | 0 | 280 | 0.49 | |
| Housing benefits (£) | 64 | 0 | 660 | 48 | 0 | 400 | 0.36 | |
| Other benefits (£) | 14 | 0 | 700 | 12 | 0 | 395 | 0.81 | |
| Observations | 295 | | | 103 | | | | |

[a]*p*-value of a t-test for equal means of the lab and online participants. [b]Average characteristics of the population of jobseeker allowance claimants in Edinburgh over the 6 months of study. The numbers are based on NOMIS statistics, conditional on unemployment duration up to 1 year. [c]High educated is defined as a university degree.

### 4.1. *Jobseeker Characteristics and Job Search History*

Demographic variables, based on the first week baseline survey and presented in Table 3, show that 43% of the lab participants are female, the average age is 36, and 43% have some university degree. Eighty percent classify themselves as 'white' and 27% have children. We can compare this to aggregate statistics about the population of jobseekers available from The Office of National Statistics (NOMIS) where we truncate unemployment duration to obtain a sample with similar median.[29] Unfortunately this provides only few variables presented in the last column of the table. It indicates that we oversample women and non-whites, while the average age is very similar. Another comparison group are the participants in our online survey which arguably face a lower hurdle to participation in the study. Results are presented in the intermediate columns, and in column 7 the *p*-value of a two-sided t-test for equal means relative to the lab participants is shown. The online survey participants differ somewhat in composition: they are more likely to be female, are slightly younger, and have less children.

The lower part of Table 3 shows variables related to job search history, also based on the first week baseline survey. The lab participants have on average applied to sixty-four jobs during the unemployment spell preceding the participation in our study. These led to 2.3 interviews and 0.42

29. Source: Office for National Statistics: NOMIS Official Labour Market Statistics. Dataset: Claimant Count conditional on unemployment duration <12 months, average over the duration of the study. Restricting attention to <12 months ensures similar median unemployment duration between the NOMIS query and our data set.

job offers.[30] Only 20% received at least one offer. Mean unemployment duration at the start of the study is 260 days, while the median is 80 days. About three-fourth of the participants had been unemployed for less than half a year. Participants typically receive jobseekers allowance and housing allowance, while the amount of other benefits received is quite low. The online survey participants are not significantly different on each of these dimensions.

To check the balance between treatment and control group we also report demographics and job search history separately by group in Table 4. Only one out of nineteen variables—the number of children—displays significant differences between the groups. This indicates balance of the sample. Balance is further corroborated by the fact that also none of the fourteen measures of search behaviour during the first 3-weeks of the study shown in the lowest panel in Table 4 displays any significant differences. We discuss these further in the next subsection. A formal assessment of balance through a Holm-Bonferroni test across all nineteen baseline variables or across all thirty-three variables including initial job search does not reject equality between the groups even at the 10% level.

### 4.2. *Descriptives of job search behaviour during the study*

In terms of job search behaviour in our study over the first three weeks, we find that the control group lists on average 493 vacancies, of which 25 are viewed, and 10 are saved (see third panel in Table 4). Out of these, participants report to have applied to 3 and eventually get an interview in 0.1 cases. Furthermore, they report about 9 weekly applications through channels outside our study, leading to 0.5 interviews on average. For the sets of listed vacancies and applications we compute a measure of occupational breadth (as described in subsection 3.7), of which the average values are also shown. Participants in the control group report 11 hours of weekly job search in addition to our study. In the weekly survey, participants were also asked to rate to what extent particular problems were a concern to them. On average, health problems are not mentioned as a major concern, while financial problems and strong competition in the labour market seem to be important. Finally, about 30% met with a case worker at the Jobcentre in a particular week.

Comparing job search behaviour and outcomes after Week 3 between treatment and control group is at the heart of the empirical assessment of the next section. Here we simply report some additional observations regarding job search behaviour to provide some background.

First, about a third of jobseekers search for jobs in the exact same occupation of their previous employment. We compare the occupations that they list in their employment history (obtained in the initial survey) with the three "preferred occupations" that they list when asked in which occupations they would prefer to find a job.[31] We find that for 35%, all of their previous occupations are now listed as preferred occupations. For 27%, some of their previous occupations are listed as preferred occupations, and for 38%, none of their previous occupations are indicated to be preferred occupations.

Secondly, most applications go to recently posted vacancies. The median age of a vacancy at the time of an application is 12 days. Of all applications to jobs from our search interface, 85% goes to a vacancy that is at most 3 weeks old at the time the application is reported.

Thirdly, the breadth and the number of vacancies that jobseekers list increase over time, while the numbers of applications and interviews decrease over time. The weekly increase in breadth

---

30. We censor the response to the survey question on the number of previous job offers at 10 and the question on interviews at 50.

31. These three preferred occupations are good proxies for actual search. When comparing them to the first occupation that is specified in the alternative interface we find that for 51% of the jobseekers this first occupation is one of the three preferred occupations.

TABLE 4
*Characteristics of the treatment and control group*

| | Control group | | | Treatment group | | | T-test |
|---|---|---|---|---|---|---|---|
| | Mean | Min | Max | Mean | Min | Max | *p*-value |
| Demographics: | | | | | | | |
| Female (%) | 42 | | | 43 | | | 0.83 |
| Age | 36 | 18 | 62 | 36 | 18 | 64 | 0.85 |
| High educ[a] (%) | 44 | | | 41 | | | 0.63 |
| Survey qualification level | 4.2 | 1 | 8 | 4.4 | 2 | 8 | 0.36 |
| White (%) | 80 | | | 80 | | | 0.97 |
| Number of children | 0.66 | 0 | 5 | 0.38 | 0 | 5 | 0.01 |
| Couple (%) | 25 | | | 21 | | | 0.41 |
| Any Children (%) | 31 | | | 24 | | | 0.17 |
| Job search history: | | | | | | | |
| Expect job within 12 weeks (%) | 0.59 | | | 0.58 | | | 0.93 |
| Vacancies applied for | 75 | 0 | 1,000 | 53 | 0 | 1,000 | 0.18 |
| Interviews attended | 2.4 | 0 | 50 | 2.2 | 0 | 50 | 0.68 |
| Jobs offered | 0.37 | 0 | 5 | 0.48 | 0 | 8 | 0.43 |
| At least one offer (%) | 20 | | | 20 | | | 0.91 |
| Days unemployed (mean) | 290 | 1 | 5,028 | 228 | 1 | 5,141 | 0.39 |
| Days unemployed (median) | 81 | 1 | 5,028 | 77 | 1 | 5,141 | |
| Less than 183 days (%) | 75 | | | 78 | | | 0.60 |
| Less than 366 days (%) | 84 | | | 87 | | | 0.54 |
| Jobseekers allowance (£) | 49 | 0 | 225 | 56 | 0 | 1,005 | 0.46 |
| Housing benefits (£) | 65 | 0 | 600 | 62 | 0 | 660 | 0.90 |
| Other benefits (£) | 9.7 | 0 | 280 | 18 | 0 | 700 | 0.41 |
| Weekly search activities in Weeks 1–3: | | | | | | | |
| Listed | 493 | 4.3 | 3,049 | 493 | 1 | 1,966 | 1.00 |
| Viewed | 25 | 3 | 86 | 26 | 0 | 119 | 0.57 |
| Saved | 10 | 0 | 65 | 11 | 0 | 79 | 0.54 |
| Applied | 3.3 | 0 | 45 | 2.5 | 0 | 33 | 0.14 |
| Interview | 0.098 | 0 | 3.3 | 0.083 | 0 | 1.5 | 0.66 |
| Applications other | 9.3 | 0 | 68 | 7.4 | 0 | 37 | 0.13 |
| Interviews other | 0.54 | 0 | 4 | 0.47 | 0 | 5 | 0.48 |
| Broadness listed[b] | 3.2 | 0 | 3.7 | 3.3 | 1 | 3.7 | 0.50 |
| Broadness applied[b] | 3 | 0 | 4 | 3.2 | 0 | 4 | 0.34 |
| Hours spent searching[c] | 11 | 0.5 | 43 | 12 | 1 | 43 | 0.15 |
| Met caseworker (%) | 0.32 | | | 0.28 | | | 0.48 |
| Observations | 152 | | | 143 | | | |

*Notes:* Demographics and job search history values are based on responses in the baseline survey from the first week of the study. Search activities are mean values of search activities over the first 3 weeks of the study.
[a]High educated is defined as a university degree.
[b]Occupational broadness, as defined in section 3.7.
[c]The number of hours spent on job search per week, as filled out in the weekly survey, averaged over Weeks 2 and 3.

of listed vacancies is about 2.2 % of a standard deviation. There is no significant trend for breadth of applications (though this is imprecisely measured), nor on weekly hours spent on job search. These results follow from regressing the outcome on a linear (weekly) time trend using only the control group and including individual fixed effects. The focus on the control group is to avoid any confounding with the treatment. The results are presented in Table 5.[32]

32. One may worry that the results in Table 5 are affected by dynamic selection as some participants leave the study over time. In Table 15 in the Online Appendix we show the results for the subsample of participants that are still present in the final weeks of the study (*i.e.* attended at least one session in Weeks 10, 11 or 12), and results are very robust.

TABLE 5
*Job search activity over time (only control group)*

|  | (1) Hours search per week | (2) Breadth of listed vac. | (3) Number of listed vac. | (4) Breadth of applications | (5) Number of applications | (6) Number of interviews |
|---|---|---|---|---|---|---|
| Time trend | 0.040 | 0.015*** | 8.91** | −0.0052 | −0.15** | −0.0072* |
|  | (0.063) | (0.0048) | (4.10) | (0.014) | (0.059) | (0.0043) |
| Individual FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Mean dep. var. | 12.2 | 3.29 | 536.1 | 3.07 | 3.38 | 0.082 |
| Weeks | 1–12 | 1–12 | 1–12 | 1–11 | 1–11 | 1–10 |
| N | 1040 | 1193 | 1196 | 504 | 1125 | 1049 |

*Notes:* All regressions contain only control group individuals. "Time trend" is a linear weekly trend. Standard errors clustered by individual in parentheses.
*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

Fourthly, we investigate whether the requirement to search on our platform has an effect on job search *per se* by comparing the patterns we just described for the control group to those for online participants who were only surveyed, including on the weekly number of job applications and interviews. Lab participants report both the numbers associated with the search in the lab and through other channels. Both groups receive no information information but one has to come physically to our lab to search on our standard interface. Lab participants report a similar number of job applications through other channels as online participants, but together with the applications associated with search in the lab the total number of applications by lab participants is significantly higher in most weeks (see Figures 10 and 11 in the Online Appendix). This difference could be the result of additional search induced through our intervention, even though we cannot rule out that it is the result of selection of more motivated participants into the lab study. In either case it clearly shows the need of a control group. The number of job interviews does not differ significantly between the groups in any week.

### 4.3. *Attrition*

The study ran for 12 weeks, but jobseekers could obviously leave the study earlier either because they found a job or for other reasons. Whenever participants dropped out, we followed up on the reasons for dropping out. In case they found a job, we asked for details, and in many cases we were able to obtain detailed information about the new job. Since job finding is a desirable outcome related to the nature of our study, we present attrition excluding job finding in Figure 1. An exit from the study is defined to occur in the week after the last session in which the individual attended a lab session. In most weeks, we lose between two and four participants, and these numbers are very similar in control and treatment groups. On average, we have 8.3 observations per participant.[33]

We now investigate whether the composition of the control and treatment group changes over time due to attrition, by looking at observable characteristics of those that remain in the study. We compute mean values of the same set of variables as in Table 4, for individuals remaining in the study in Weeks 1, 4, and 12. For each of these groups of survivors, we test whether the treatment and control group are significantly different. Since we present thirty-two variables for three groups of survivors, this implies ninety-six tests. The resulting *p*-values are presented in

---

33. In the Online appendix we show the distribution of the number of attended weeks per participant, split by pre-intervention (Weeks 1–3) and post-intervention (Weeks 4–12). See Figures 7–9.
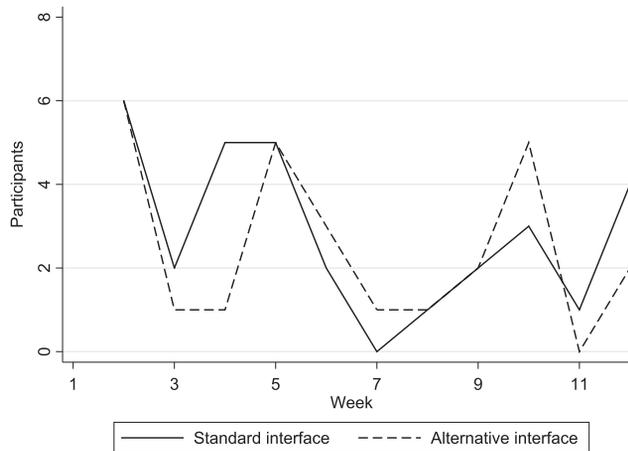
FIGURE 1

Attrition of participants in the standard and alternative interface groups (excluding job finding).

Table 31 in the Online Appendix. Only six of the *p*-values are smaller than 0.10, so there is no indication that attrition leads to systematic differences in the composition of the treatment and control group. Also a Holm–Bonferroni test for joint significance does not reject the null hypothesis of identical values.

The apparent lack of selection is on the one hand helpful to study how the intervention may have affected search outcomes, on the other hand it hints that we are unlikely to capture differences in job finding rates, which are low overall. We will come back to the analysis of drop out and job finding in more detail in subsection 5.5.

### 4.4. *Use of alternative interface*

An obvious question regarding our treatment intervention is whether participants actually use the alternative interface. They are free to revert back to the standard interface, and in this sense our intervention can be considered an intention-to-treat. We are hesitant to adopt this interpretation since all participants in the treatment group used the alternative interface at least once and were therefore exposed to recommendations and suggestions based on their declared "desired" occupation. It could be that they used this information when they reverted back to searching with the standard interface. With this in mind, we report information on actual usage. Panel (a) of Figure 2 plots the fraction of users of the alternative interface over the 12 weeks. On average we find that around 50% of the listed vacancies of the treated participants come from searches using the alternative interface over the 8 weeks and this fraction remains quite stable throughout.[34] We discuss panel (b) that considers subgroups of participants later on.

---

34. The variation in usage results from both between and within users. The participants in the treatment group use the alternative interface for at least one search in 75% of the weeks on average, and in 35% of the weeks the alternative interface was used solely. These findings are shown in Figure 13 (Panel (a)) in the Online Appendix. Further statistics on usage are shown in Panels (b) and (c).
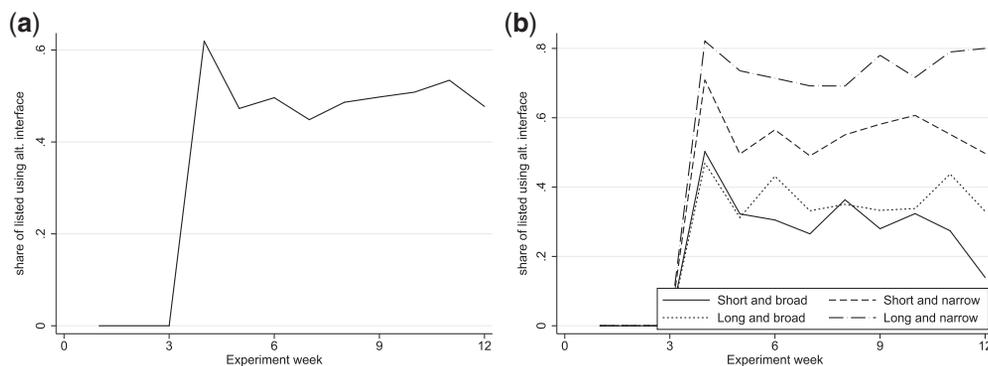
FIGURE 2

Share of listed vacancies that results from using the alternative interface. (a) Average in treatment group. (b) Average in treatment group by type (unemployment duration and occupational breadth).

## 5. ANALYSIS AND RESULTS

As outlined in the Introduction, the hypothesis behind the intervention is that providing information about other occupations will allow individuals to explore vacancies from a larger set of occupations, which may lead to more job interviews. This should hold in particular for individuals that otherwise explore few occupations. Exploring more occupations could go along with more search, or with the same search effort concentrated on more occupations but in a closer geographic region. The following lays out the empirical strategy to investigate this.

### 5.1. *Econometric specification*

Our data are a panel and our unit of observation is at the week/individual level. That is, we compute a summary statistic for each individual of her search behaviour (vacancies listed, applications, interviews) in a given week; see Section 3.7 for a description of the outcome measures of interest. Since it is a randomized controlled experiment in which we observe individuals for 3 weeks before the treatment starts, the natural econometric specification is a model of difference-in-differences. To take account of the panel structure we include individual random effects. By design, there should be no correlation between individual characteristics (observable and unobservable) and treatment assignment, at least initially. To test whether the random effects specification is appropriate for the entire duration of the study, we have estimated a fixed effects model and performed a Hausman test for each of the main specifications (see Table 17 in Online Appendix 8.3). In none of the cases we could reject that the random effects model is consistent, such that we decide in favour of the random effects model for increased precision. We discuss robustness at the end of this section (subsection 5.6) and show that point estimates are similar when using individual fixed effects yet precision is lower. As has been emphasized by Bertrand *et al.* (2004), serial correlation is an issue in difference-in-differences models. We follow their suggestion and average the weekly observations into two observations per individual, one before (Weeks 1–3) and one after the intervention (Weeks 4–12), but again report robustness to alternative specifications at the end of this section.

We compare a variable measuring an outcome ($Y$) in the control and treatment group before and after the week of intervention, controlling for time period fixed effects ($\alpha_t$, before or after the intervention), time–slot × wave fixed effects ($\delta_g$) and a set of baseline individual characteristics ($X_i$) to increase the precision of the estimates. The treatment effect is captured by a dummy variable

$(T_{it})$, equal to 1 for the treatment group in the period after the intervention. The specification is:

$$Y_{it} = \alpha_t + \delta_g + \gamma T_{it} + X_i \beta + \eta_i + \epsilon_{it} \tag{1}$$

where $i$ relates to the individual, $t$ to the time period and $\eta_i + \epsilon_{it}$ is an error term consisting of an individual specific component ($\eta_i$) and a white noise error term ($\epsilon_{it}$). Individual characteristics $X_i$ include gender, age and age squared, unemployment duration and unemployment duration squared and dummies indicating financial concerns, being married or cohabiting, having children, being highly educated and being white. Standard errors are clustered at the individual level in the regressions, to account for any remaining correlation of an individual's observations.

As mentioned earlier, one important challenge with such approach has to do with attrition. If there is differential attrition between treatment and control groups, it could be that both groups differ in unobservables following the treatment. We proceed in two ways to address this potential concern. First, in Section 4.3, we document attrition across treatment and control groups and find no evidence of asymmetric attrition in terms of observable characteristics. Secondly, our panel structure allows us to control for time-invariant heterogeneity and use within-individual variation. When we estimate a random and fixed effects model, as mentioned above the Hausman test fails to reject the latter. Even though the treatment itself is assigned at the group-level and it is unlikely to be correlated with unobserved individual characteristics, differential attrition could create correlation between the treatment and unobservable individual characteristics. This would then lead to rejection of the random-effects model. The fact that we can never reject this model is thus another indication against differential attrition between treatment and control groups.

It is likely that the treatment affects different individuals differentially. In order for our intervention to affect job search and job prospects, it has to open new search opportunities to participants and participants have to be willing to pursue those opportunities. Participants may differ in terms of their search strategies. We expect our intervention to broaden the search for those participants who otherwise search narrowly, which we will measure by their search in the weeks prior to the intervention. For those who are already searching broadly in the absence of our intervention it is not clear whether we increase the breadth of their search. We therefore estimate heterogeneous treatment effects by initial breadth (splitting the sample at the median level of breadth over the first 3 weeks).[35]

Secondly, the willingness to pursue new options depends on the incentives for job search, which change with unemployment duration for a variety of reasons. Longer-term unemployed might be those for whom the search for their preferred jobs turned out to be unsuccessful and who need to pursue new avenues, while they are also exposed to institutional incentives to broaden their search (the Jobcentres require jobseekers to become broader after three months). We therefore also interact the treatment effect with a dummy for above median unemployment duration.[36]

35. To check the robustness of our classification of jobseekers as narrow or broad searchers, we computed three different classifications (based on listed vacancies in Week 1, Week 2, and Week 3). We find that the classifications of Weeks 1 and 2 agree on 69% of the jobseekers, those of Weeks 1 and 3 agree on 67% of the jobseekers, and those of Weeks 2 and 3 agree on 86% of the jobseekers.

36. When estimating heterogeneous effects we adapt our specification to include all necessary additional terms. Define $D_i$ to be an indicator equal to one for individuals belonging to group 1 (*e.g.* narrow searchers) and equal to zero for individuals belonging to group 2 (*e.g.* broad searchers). We estimate:

$$Y_{it} = \theta D_i + \alpha_{1t} D_i + \alpha_{2t}(1-D_i) + \delta_g + \gamma_1 T_{it} D_i + \gamma_2 T_{it}(1-D_i) + X_i \beta + \eta_i + \epsilon_{it} \tag{2}$$

Thus, the specification contains an additional baseline difference between the groups ($\theta$), differential time period effects for the the groups ($\alpha_{1t}$ and $\alpha_{2t}$) and differential treatment effects between the groups ($\gamma_1$ and $\gamma_2$). Note that since we average observations into two periods (before and after the intervention), $\alpha_{1t}$ and $\alpha_{2t}$ simply contain a time effect for the

Apart from these dimensions for which we have clear reasons for separate investigation we do not explore other dimensions of heterogeneity to avoid data mining. Nevertheless, it might be interesting to know whether breadth of search is correlated with other factors that might drive the observations we report. We investigate this by regressing it on a number of individual characteristics. Results are presented in columns (1) and (2) in Table 16 in the Online Appendix. We find that breadth of search is not easily predicted based on individual characteristics. Almost all variables are not statistically different from zero, and the $R^2$ of the regression is low (0.18). The same holds for unemployment duration (columns (3) and (4)).

For the sake of brevity, in the main body we only present the results on the treatment effect ($\gamma$) as well as the interaction effects between the treatment and the subgroups of interest. In Table 22 in the Online Appendix, we report full results including all other covariates for the main regressions.

### 5.2.  *Effects on listed vacancies*

We first look at the effects on listed vacancies—both in terms of number and breadth. We have two variables measuring how broad participants search, one in terms of occupation (as described in Section 3.7), the other in terms of geography (fraction of vacancies outside Edinburgh metropolitan area).

We estimate equation (1) and present results in Table 6. The first row shows a significant positive overall effect on breadth of search in terms of occupation. The breadth measure increases with 0.13, which amounts to approximately one-fifth of a standard deviation. Another way to assess the magnitude of this effect is to compare it to the natural increase in breadth of listings over time (as shown in Table 5), which implies that the treatment effect is equivalent to the broadening that on average happens over 9 weeks. We find no significant evidence of an overall effect on geographical breadth or on the number of listed vacancies.

In rows (2) and (3) in Table 6, we split the sample according to how occupationally broad jobseekers searched in the first 3 weeks. We find clear heterogeneous effects: those who looked at a more narrow set of occupations in the first 3 weeks become broader, while those who were broad become more narrow as a result of the intervention. Note that these effects are not driven by "regression to the mean" since we compare narrow/broad searchers in our treatment group to similarly narrow/broad searchers in our control group.[37] We again find no significant effects on the geographic distance of job search nor on the number of job applications.[38,39] The total effects on job prospects remains in either case an empirical matter that we take up in subsequent sections. The different effects on occupational breadth can be reconciled in a setting where broad searchers

---

second period. Note also that, just as in the baseline model, the specification contains time slot X wave dummies ($\delta_g$) and since treatment is assigned at the time-slot-level, these control for any baseline differences between the control and treatment group.

   37.  In Figure 12 in the Online Appendix, we show the mean breadth of the different groups before and after the intervention to clarify further that these results are not caused by regression to the mean.

   38.  In the Online Appendix we also report estimates where we split the sample according to breadth along the geographical dimension at the median (see Table 19). The results are similar (those who were searching broadly become more narrow and vice versa, and there is some trade-off with occupational breadth). This could still be driven by initial occupational breadth, since this is negatively correlated with initial geographical breadth (coefficient $-0.36$) and is not controlled for. Indeed, when we split both by occupational and geographical breadth the effects are driven by the occupational dimension, which we will henceforth focus on.

   39.  The difference in the number of observations between the columns in Table 6 and similar tables that follow is due to the fact that we can only compute the occupational (geographical) breadth measure if the number of listed is two (one) or larger, which excludes different numbers of observations depending on the variable of interest.

TABLE 6
*Effect of intervention on listed vacancies*

| | Breadth of listings | | Number of listings |
|---|---|---|---|
| | (1) Occupational | (2) Geographical | (3) Lab |
| Treatment | 0.13** | −0.01 | −34.99 |
| | (0.06) | (0.02) | (52.09) |
| Treatment | | | |
| X occupationally broad | −0.07** | 0.01 | −23.71 |
| | (0.04) | (0.03) | (90.08) |
| X occupationally narrow | 0.34*** | −0.03 | −41.84 |
| | (0.10) | (0.03) | (64.01) |
| Model | Linear | Linear | Linear |
| Observation weeks | 1–12 | 1–12 | 1–12 |
| N | 540 | 541 | 541 |

*Notes:* Each column represents two separate regressions. All regressions include time-slot fixed effects, period fixed effects (separately for each subgroup), individual random effects and individual characteristics. Standard errors clustered by individual in parentheses.
**$p < 0.05$, ***$p < 0.01$.

TABLE 7
*Effect of intervention on listed vacancies—interactions*

| | Breadth of listings | | Number of listings |
|---|---|---|---|
| | (1) Occupational | (2) Geographical | (3) Lab |
| Treatment | | | |
| X long unempl. and occ. broad | −0.10** | 0.06 | 189.12 |
| | (0.05) | (0.04) | (135.01) |
| X short unempl. and occ. broad | −0.05 | −0.04 | −252.80** |
| | (0.05) | (0.05) | (120.19) |
| X long unempl. and occ. narrow | 0.36** | −0.04 | 23.35 |
| | (0.15) | (0.05) | (62.51) |
| X short unempl. and occ. narrow | 0.32** | −0.01 | −112.82 |
| | (0.13) | (0.05) | (116.52) |
| Model | Linear | Linear | Linear |
| Observation weeks | 1–12 | 1–12 | 1–12 |
| (N) | 540 | 541 | 541 |

*Notes*: Each column represents one regression. All regressions include time-slot fixed effects, period fixed effects (separately for each subgroup), individual random effects and individual characteristics. Standard errors clustered by individual in parentheses.
**$p < 0.05$.

find many occupations plausible and use the additional information to narrow down the suitable set, while narrow searchers find few occupations suitable and use the additional information to broaden this set. This mechanism is more formally described in Section 6.

Finally, we split the effect further depending on how long jobseekers have been searching for a job and present the results in Table 7. We interact the intervention effect with two groups: short-term unemployed (with unemployment duration of less than the median of 80 days) and long-term unemployed (with unemployment duration above the median). We find that results do not change much, though standard errors are larger. We still find that occupationally narrow

TABLE 8
*Effect of intervention on applications*

| | Breadth of applications | | Number of applications | | |
|---|---|---|---|---|---|
| | (1) Occupational | (2) Geographical | (3) Lab | (4) Outside lab | (5) Total |
| Treatment | 0.03 | −0.06* | 0.09 | −0.03 | 0.01 |
| | (0.20) | (0.03) | (0.16) | (0.09) | (0.09) |
| Treatment | | | | | |
| X occupationally broad | −0.43* | −0.02 | −0.08 | −0.06 | −0.05 |
| | (0.22) | (0.05) | (0.19) | (0.13) | (0.12) |
| X occupationally narrow | 0.49* | −0.09** | 0.27 | −0.02 | 0.08 |
| | (0.29) | (0.04) | (0.27) | (0.13) | (0.13) |
| Model | Linear | Linear | Neg. Bin. | Neg. Bin. | Neg. Bin. |
| Observation weeks | 1–11 | 1–11 | 1–11 | 1–11 | 1–11 |
| (*N*) | 305 | 363 | 541 | 490 | 487 |

*Notes*: Each column represents two separate regressions. All regressions include time-slot fixed effects, period fixed effects (separately for each subgroup), individual random effects and individual characteristics. Columns (3)–(5) are Negative Binomial regression models where we report [exp(coefficient) − 1], which is the percentage effect. Standard errors in parentheses (clustered by individual in columns (1) and (2)).
*$p < 0.10$, **$p < 0.05$.

searchers become broader while those that were already broad become more narrow, irrespective of unemployment duration. Shorter unemployed jobseekers seem to consult less listings in the treatment group, significantly so for broader ones. If the new information allowed them to focus their search better this might not necessarily harm their job prospects, as outlined in our theoretical model, but nevertheless this remains a concern that we return to when we consider the effect on job interviews.

### 5.3. *Effects on applications*

The second measure of search behaviour relates to applications. We have information about applications based on search activity conducted inside the laboratory as well as outside the laboratory. The distribution of applications contains a large share of zeroes (in almost 50% of the weekly observations there are zero applications through the lab). Therefore, we estimate a negative binomial model, with individual random effects.[40] For these models we report [exp(coefficient) − 1], which is the percentage effect.

The results are presented in Table 8. We find no overall treatment effect on applications, except for a decrease in their geographical breadth (approximately one-fifth of a standard deviation). When we split the sample according to initial occupational breadth, we find a similar pattern as for listings. Those who searched more narrowly in terms of occupation become occupationally broader, while those that searched broadly become more narrow. The estimates are significantly

40. Due to overdispersion in the distribution of applications, we prefer a negative binomial model over a Poisson model. However, negative binomial regressions are sometimes less robust and in addition no consensus exists on how to include fixed effects (Allison and Waterman, 2002). Furthermore, we can not cluster standard errors with the random effects negative binomial regressions. Therefore we also report results from Poisson regressions in Online Appendix 8.3 (Table 18). The findings are similar. Furthermore, as we average weekly observations into a before and after period, the outcome variable is changed from discrete into continuous. The distribution still resembles a Poisson distribution and a Poisson regression model or a negative binomial model can still be used (see Gourieroux *et al*., 1984).

TABLE 9

*Effect of intervention on applications—interactions*

|  | Breadth of applications | | Number of applications | | |
|---|---|---|---|---|---|
|  | (1) Occupational | (2) Geographical | (3) Lab | (4) Outside lab | (5) Total |
| **Treatment** |  |  |  |  |  |
| X long unempl. and occ. broad | −0.67*** | −0.07 | −0.24 | −0.22 | −0.20 |
|  | (0.25) | (0.06) | (0.21) | (0.14) | (0.13) |
| X short unempl. and occ. broad | −0.18 | 0.02 | 0.17 | 0.17 | 0.17 |
|  | (0.33) | (0.07) | (0.36) | (0.22) | (0.21) |
| X long unempl. and occ. narrow | 0.51 | −0.10** | 0.42 | −0.11 | 0.00 |
|  | (0.34) | (0.05) | (0.40) | (0.16) | (0.17) |
| X short unempl. and occ. narrow | 0.40 | −0.08 | 0.25 | 0.14 | 0.22 |
|  | (0.41) | (0.06) | (0.37) | (0.20) | (0.21) |
| Model | Linear | Linear | Neg. Bin. | Neg. Bin. | Neg. Bin. |
| Observation weeks | 1–11 | 1–11 | 1–11 | 1–11 | 1–11 |
| *N* | 305 | 363 | 541 | 490 | 487 |

*Notes*: Each column represents one regression. All regressions include time-slot fixed effects, period fixed effects (separately for each subgroup), individual random effects and individual characteristics. Columns (3)–(5) are Negative Binomial regression models where we report [exp(coefficient)−1], which is the percentage effect. Standard errors in parentheses (clustered by individual in columns (1) and (2)).
**$p < 0.05$, ***$p < 0.01$.

different from zero at the 10% level. We find no effects on the number of applications for either group (columns (3)–(5)), though point estimates might indicate a pattern where the initially-narrow group expands broadness through more applications and vice versa for the initially-broad subgroup. There is also a negative effect on geographical breadth for the occupationally narrow jobseekers (column (2)), which indicates that narrow jobseekers search for occupationally broader jobs closer to home.[41]

Again, we split these effects by the duration of unemployment. Column (1) in Table 9 shows that occupational breadth goes down significantly for long-term unemployed broad searchers. It increases most for long-term unemployed narrow searchers, yet this is insignificant due to large standard errors. This increase is accompanied by a significant decrease in geographical distance. Estimates on the number of applications are insignificant, though point estimates are economically large. As noted earlier, even decreases in occupational breadth can be beneficial if job search becomes better targeted.

### 5.4. *Effects on interviews*

We now turn to interviews, the variable that is most closely related to job prospects. Since the number of interviews per week is always very small, we cannot compute breadth measures. So we only look at a measure of the number of interviews obtained as a result of search conducted inside the laboratory and outside the laboratory.[42] Because of the large share of zeros, we estimate

41. When splitting the sample according to how narrow people searched in terms of geography, we find no evidence of heterogeneous effects. Results are presented in the Online Appendix in Table 20.

42. For interviews reported outside the lab we censor observations at three interviews per week, because of some outliers. Results are similar when no such restriction is imposed. As a check of consistency, we also check whether interviews are ever reported without preceding applications. We find that in 98.2% of the weeks in which an interview is reported, a positive number of applications was reported in at least one of the two preceding weeks.

TABLE 10
*Effect of intervention on interviews*

|  | Number of interviews | | |
|---|---|---|---|
|  | (1) Lab | (2) Survey | (3) Total |
| Treatment | 0.61 | 0.40* | 0.44* |
|  | (0.79) | (0.27) | (0.28) |
| Treatment | | | |
| X occupationally broad | −0.37 | −0.00 | −0.07 |
|  | (0.43) | (0.28) | (0.24) |
| X occupationally narrow | 1.13 | 0.86** | 1.03*** |
|  | (1.26) | (0.47) | (0.55) |
| Model | Poisson | Poisson | Poisson |
| Observation weeks | 1–10 | 1–10 | 1–10 |
| $N$ | 540 | 466 | 464 |

*Notes*: Each column represents two separate regressions. All regressions include time-slot fixed effects, period fixed effects (separately for each subgroup), individual random effects and individual characteristics. Columns (1)–(3) are Poisson regression models where we report [exp(coefficient) − 1], which is the percentage effect. Standard errors clustered by individual in parentheses.
*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

a Poisson model with individual random effects. Again we report [exp(coefficient) − 1], which is the percentage effect.

Results are presented in Table 10. There is a positive effect of the treatment of 44% on the total number of interviews, which is significant at the 10% level. We also find positive effects on interviews on the two separate dimensions of search in the lab and search outside the lab, but even though the point estimate for the effect within the lab is highest only the increase in out-of-lab interviews is statistically significant. This can be explained by the difference in base rate which is lower in the lab making statistical inference more difficult: In the pre-treatment period the number of interviews through the lab was 0.09, while the number of interviews through other channels was 0.53.

When splitting the sample according to breadth of search, we find that the effect is entirely driven by those who searched narrowly in terms of occupation. For this group the number of interviews increases for search activity conducted both in the lab and outside (though again, only the increase of the out-of-lab interviews is statistically significant). This seems to indicate that the additional information is not only helpful for search on our platform, but also guides behaviour outside.[43] The point estimates for the occupationally broad group are all insignificant and in absolute value much smaller, but point estimates are negative.

When we further split the sample according to the length of unemployment duration, we find that the positive treatment effects on the narrow searchers is mainly driven by the long-term unemployed narrow searchers (Table 11). This group gets a significant increase in the number of interviews both as a result of search activity inside the lab and outside the lab.[44] This highlights

43. We find some evidence of heterogeneity in treatment effects when we split the sample according to initial geographical breadth, with a large positive significant treatment effect for those who searched broadly geographically. Results are presented in the Online Appendix in Table 21.

44. The extremely large value of the increase in lab interviews for the long term narrow searchers is partly due an individual outlier that reported an average of 3.5 interviews per week in the treatment period. If we exclude this individual, the coefficient is still large and positive (6.75, significant at the 1% level).

TABLE 11
*Effect of intervention on interviews—interactions*

| | Number of interviews | | |
|---|---|---|---|
| | (1) Lab | (2) Survey | (3) Total |
| Treatment | | | |
| X long unempl. and occ. broad | −0.27 | −0.23 | −0.21 |
| | (0.72) | (0.25) | (0.26) |
| X short unempl. and occ. broad | −0.37 | 0.17 | 0.01 |
| | (0.52) | (0.47) | (0.37) |
| X long unempl. and occ. narrow | 13.12*** | 2.44*** | 3.39*** |
| | (9.25) | (1.19) | (1.42) |
| X short unempl. and occ. narrow | −0.26 | 0.31 | 0.30 |
| | (0.51) | (0.40) | (0.44) |
| Model | Poisson | Poisson | Poisson |
| Observation weeks | 1–10 | 1–10 | 1–10 |
| *N* | 540 | 466 | 464 |

*Notes*: Each column represents one regression. All regressions include time-slot fixed effects, period fixed effects (separately for each subgroup), individual random effects and individual characteristics. Columns (1)–(3) are Poisson model regressions where we report $[\exp(\text{coefficient}) - 1]$, which is the percentage effect. Standard errors clustered by individual in parentheses. ***$p < 0.01$.

TABLE 12
*Effect of intervention on interviews: core and non-core occupations*

| | Number of interviews (in the lab) | |
|---|---|---|
| | (1) Core occupations | (2) Non-core occupations |
| Treatment | −0.14 | 0.75 |
| | (0.72) | (0.85) |
| Model | Poisson | Poisson |
| Observation weeks | 1–10 | 1–10 |
| *N* | 540 | 540 |

*Notes*: Each column represents one regression. All regressions include time-slot fixed effects, period fixed effects (separately for each subgroup), individual random effects and individual characteristics. Columns (1)–(2) are Poisson model regressions where we report $[\exp(\text{coefficient}) - 1]$, which is the percentage effect. Standard errors clustered by individual in parentheses.

that our intervention is particularly beneficial to people who otherwise search narrowly and who have been unemployed for some months. It might be encouraging that there are no significant negative effects on the groups that became occupationally narrower, but some negative point estimates might warrant further investigation.

The set of weekly interviews is too small to compute breadth measures. We did, however, ask individuals at the beginning of the study to indicate three core occupations in which they search for jobs, and we observe whether an interview was for a job in someone's core occupation or for a job in a different occupation. We had seen earlier that the alternative interface was successful in increasing the occupational breadth of listed vacancies and applications, and separate treatment effects on interviews in core versus non-core occupations allow some assessment of whether this lead to more "breadth" in job interviews. Results are presented in Table 12. We indeed find that

the increase in the number of interviews relative to the control group comes from an increase in non-core occupations that were not their main search target at the beginning of our study, though due to low precision the effect is not statistically significant. As the number of interviews becomes small when splitting between core and non-core, we cannot split the sample further by subgroups.

One may worry that the increase in interviews in non-core occupations is associated with different quality of the interviews. For example, the suggestions could lead to interviews for jobs with different wages. We have investigated this by comparing the average wage of listed vacancies, applications, and interviews and find that the alternative interface does not significantly change the wage of any of these.[45]

Our findings suggest that the alternative interface may be more beneficial to those that search narrowly and have been relatively long unemployed. This finding is supported by statistics on usage of the interface over time. Panel (b) of Figure 2 shows the evolution of the fraction of treated participants using the interface, splitting the sample by occupational breadth and unemployment duration. We find that long-term narrow searchers are indeed using the interface more than the other groups (with around 75% of them using the interface in contrast to around 45% for the other groups), and this difference is statistically significant. This finding supports the intuition that some groups of jobseekers benefit more from the intervention and are therefore more willing to use the alternative interface. This group, the long-term unemployed narrow searchers is exactly the group for which we find the most pronounced positive effects.[46]

### 5.5. *Effects on job finding*

We now briefly turn to the analysis of job finding. As mentioned earlier, the study was not designed to evaluate effects on job finding and, given the size of the sample, we should be cautious in interpreting any results we have. Also, one should keep in mind that attrition from one week to the next for unexplained reasons is low but of the same order of magnitude as the confirmed job finding rate.[47]

At the end of the 12 weeks, 28% of the participants using the standard interface have found a job, against 22% of the participants using the alternative interface. A similar proportion (15%) of participants have dropped out of the study with an unclear outcome, so it is difficult to draw conclusions based on these numbers.

These numbers are nevertheless useful to get a sense of the sample size one would need in order to capture significant effects on job finding. We perform a simple sample size calculation to illustrate how the required sample size for finding an effect on job finding exceeds the sample size required for finding an effect on the number of interviews. To detect a 44% increase in interviews due to the intervention (see Table 10), a sample size of 70 observations per treatment is required (so 140 in total). For job finding, detecting a similar sized effect requires around 3,794 observations per treatment, due to a much lower base rate.[48] Even if one takes the (at most) eight

---

45. We computed for every individual in every week the average wage of listed vacancies, applications or interviews and performed regressions similar to our main specifications.

46. The idea that these groups are more willing to use the alternative interface is supported by responses from the baseline survey in the first week. The participants were asked to specify how long they expected it would take to find a job. Within the group of short-term unemployed the median response is "less than 3 months" which might indicate a rather clear idea of how to obtain a job, while for the long-term unemployed group the median response is "less than 6 months" which might indicate a less clear view and more scope to provide successful alternatives.

47. We tried to follow-up by calling them at least three times, though for a non-trivial share of the attrition we still do not observe whether the person found a job or just quit the study.

48. The precise computation is as follows. We observe in the first three weeks that, on average, participants have a total of 0.61 interviews per week through the lab and other channels (see Table 4). To detect a 44% increase in interviews

TABLE 13
*Treatment effects on job finding rate*

|  | (1) | (2) |
|---|---|---|
| Treatment | −0.13 | −0.18 |
|  | (0.25) | (0.31) |
| Treatment x Occupationally narrow |  | 0.10 |
|  |  | (0.56) |
| N | 253 | 253 |

*Notes*: Proportional Cox Hazard model, with time-slot fixed effects, and individual characteristics. We exclude observations censored at 3 weeks or less. Reported values are coefficients.

observations per individual in our study into account, it is clear that we lack power to identify any realistic effect on job finding.

Bearing this in mind, we estimate a simple duration model where the duration is the number of weeks we observe an individual until she/he finds a job. Since we know when each individual became unemployed, we can calculate the total unemployment duration and use this as a dependent variable. This variable is censored for individuals who drop out of the study or who fail to find a job before the end of the study. We estimate a proportional Cox hazard model with the treatment dummy as independent variable, controlling for additional individual characteristics and group session dummies.

We report estimates for the entire sample and for the sub-samples conditioning on initial search type (narrow versus broad search). The results are presented in Table 13. We fail to find significant differences in the hazard rates across treatments. That is, we have no evidence that the jobseekers exposed to the alternative interface were more or less likely to find a job (conditionally on still being present in Week 4). Despite the negative point estimate for the treatment group, even increases in the hazard of the treatment group of the magnitude of the increase in interviews overall (29%) or for narrow individuals (52%) are well within the confidence interval of these estimates. That is not to say that lack of power is the only plausible reason for finding no effect. As mentioned in the introduction, interviews in broader occupations might not convert to jobs at the same rate. We return to advocating larger studies in the conclusion.

### 5.6. *Robustness: alternative specifications*

In our analysis we made some choices regarding the empirical specification and the definition of variables. Below we briefly discuss alternative choices and investigate robustness of our results (more details can be found in Online Appendix 8.4). We consider (1) individual fixed effects instead of random effects, (2) weekly observations instead of aggregated data, (3) linear models instead of count data models, (4) excluding the last one or two observations per individual, (5)

due to the intervention (see Table 10), such that the interview rate becomes 0.89, a sample size of 70 observations per treatment is required (so 140 in total). This number is based on an one-sided test with type-I error probability $\alpha = 0.10$ and power $1 - \beta = 0.80$. The standard deviation is assumed to be 0.75 in both groups, based on the numbers reported in Table 4. For job finding, we observe ninetten people finding a job in the first 3 weeks, which implies a weekly job finding rate of approximately 0.02. If we make the (strong) assumption that the additional interviews are equally likely to result in a job as the initial interviews, we would expect a 44% increase in job finding. Note that this is a conservative choice as this would be a very large effect. Still, to be able to pick up the increase in job finding from 0.02 to 0.0288 requires a sample size of 3,794 people per treatment (similar test as for interviews).

an alternative breadth measure, and (6) IV regressions with the use of the alternative interface as the treatment intensity.[49]

Our specifications include individual random effects to increase precision. A Hausman test does not reject validity of the random effects model. In Table 23 of the Online Appendix we show our baseline regressions using individual fixed effects instead of random effects. We find very similar overall patterns but reduced precision and significance. In particular, changes in occupational breadth are similar (for listed vacancies and applications), while we find large positive coefficients for interviews for narrow searchers, which are, due to slightly reduced precision not statistically significant.

All data in our estimations have been averaged into two periods, before and after the intervention (as suggested by Bertrand *et al.* (2004)). In Table 24 in Online Appendix 8.3, we show that results are very similar when including weekly observations, both for changes in breadth and for the number of interviews.

Since the number of applications and interviews are count variables, we use Poisson regressions or negative binomial regressions in our analysis. In Table 25 in Online Appendix 8.3, we present linear regressions for these outcomes. We find similar patterns: there is no clear impact on applications, but the point estimate for interviews is economically large, and significant for narrow searchers.

In our analysis we excluded Week 12 (for applications) and Weeks 11 and 12 (for interviews), because for vacancies saved in these weeks we can not follow up on whether an application was sent or an interview was secured. Alternatively, we can exclude *for each individual* their final one or two attended sessions. The results when using this approach are shown in the Online Appendix in Table 26. All findings are very similar, both in magnitude and statistical significance.

Fifthly, we consider our definition of occupational breadth. In our approach the distance between two occupations is based on the number of common digits of the two occupational codes (see Subsection 3.7). We consider two alternatives. First, one can call occupations identical if they share a particular number of digits, which leads to the well-known Gini–Simpson index. Such measures are highly correlated with our measure (for listed vacancies our measure has a correlation above 0.95 with four different Gini–Simpson measures, see Tables 29 and 30 in the Online Appendix). Not surprisingly we find very similar results when we adopt this alternative measure. Secondly, we could use a more elaborate alternative which is to use empirically observed transitions between occupations in labour market surveys to measure the "closeness" of the two occupations.[50] The results of the main regressions using this approach are presented in Table 27 in Online Appendix 8.3. We find, again, that results are very robust against this alternative breadth measure.

Finally we consider an interpretation that all our results are intention-to-treat effects (due to voluntary usage of the alternative interface). As discussed, we are hesitant to emphasize this interpretation too much, because suggestions about alternative occupations can affect jobseekers even after a user switches back to the standard interface. However, for the sake of comparison, we can consider treatment assignment as an instrument for actual usage when estimating our empirical models.[51] The results of such an approach are presented in Table 28. As expected, the estimates are larger in magnitude. We find that breadth of listed vacancies increases (with

---

49. We also thank an anonymous referee for requesting additional analysis of heterogeneous effects by educational level. We find no evidence. In this article we just focus on two obvious dimensions of heterogeneity, to prevent data mining.

50. We thank an anonymous referee for this suggestion.

51. We define actual usage as the share of listings that is performed using the alternative interface in a particular week.

a coefficient of 0.24\*\* compared to 0.13\*\* in the baseline results). Additionally, we also find that the breadth of applications increases significantly for narrow searchers. Interviews increase significantly for narrow searchers.

## 6. AN ILLUSTRATIVE MODEL

In the empirical section we saw that our information intervention increases occupational breadth: listings are broader and more job interviews are obtained, possibly driven by jobs outside the core occupations. Effects are concentrated on those who initially search narrowly, while breadth decreases for those who initially already search broadly. Finally and possibly least obvious, effects are strongest for the longer-term unemployed. Here we briefly sketch a very stylized occupational job search model that is capable of organizing our thoughts about the driving forces. It is based on the idea that workers learn about the occupations in which they search for jobs, in the spirit of for example, Neal (1999), with the difference that workers start with heterogeneous beliefs about different occupations and that we study information provision.

A jobseeker can search for jobs in different occupations, indexed $i \in \{1,\ldots,I\}$. For each occupation she decides on the level of search effort $e_i$. Returns to searching in occupation $i$ are given by an increasing but concave function $f(e_i)$.[52] The returns to getting a job are given by wage $w$ and are the same across occupation, and $b$ denotes unemployment benefits. The cost of search is given by an increasing and convex function $c(\sum e_i)$. A limiting case is a fixed total search effort $\bar{e}$, such that costs are zero up to that point and infinite thereafter.

The individual is not sure of her job prospects within the various occupations. In a given occupation $i$ her job prospects are either good (in which case we denote her $H$—high—type) and she obtains a job with arrival probability $a_H f(e_i)$. Otherwise her prospects are bad (in which case we denote her $L$—low—type) and her job probability is $a_L f(e_i)$, where $a_H > a_L = 0$, where the equality is only for simplicity. The individual does not know whether she is a high or low type in occupation $i$, but assigns probability $p_i$ to being a high type. So the individual's belief is a vector $(p_1, p_2, \ldots, p_I)$ of probabilities for each occupation. Types are drawn iid and therefore the type vector is all that is relevant for the individual. Still, to model the information content of the alternative interface later on, it will be convenient to make the additional assumption that the individual is unsure of the exact value of the probability in each of the occupations, and only knows its distribution $Q_i$ with support $[\underline{q}_i, \overline{q}_i]$ among people that are like her. Then $p_i$ can be interpreted as the average belief according to $Q_i$. For technical convenience assume that types are not too good, that is, $\overline{q}_i \leq 1/2$, so that the average belief is also bounded by this number. This ensures that an occupation with higher belief also has higher variance and both increase the incentives to search in this occupation in such a simple bandit problem, which makes search incentives monotone in $p_i$.

Given a belief vector $p = (p_1, \ldots, p_I)$ and a vector of search effort in the various occupations $e = (e_1, \ldots, e_I)$, her overall expected probability of being hired is

$$H(p,e) = 1 - \prod_i [1 - f(e_i)(p_i a_H + (1 - p_i) a_L)]$$

where the product gives the probability of not getting a job offer in any occupation.

Assume the unemployed jobseeker lives for $T$ periods, discounts the future with factor $\delta$, and there are no job separations. An individual who has a prior $p_i^t$ about his type in occupation $i$ at

---

52. This could embed a standard matching function with search effort—we do not go into the foundations here.

the beginning of period $t$ and spends effort $e_i^t$ during the period but does not get a job will update her beliefs by Bayes rule. Let $B(p_i^t,e_i^t)$ denote this new belief. For interior beliefs we have[53]

$$p_i^{t+1} = B(p_i^t,e_i^t) = \begin{cases} = p_i^t & if \ \ e_i^t = 0 \\ < p_i^t & if \ \ e_i^t > 0, \end{cases} \qquad (3)$$

since the individual becomes more pessimistic if unsuccessful unless she put zero effort in the first place. Let $B(p,e) = (B(p_1,e_1),\ldots,B(p_I,e_I))$ denote the vector of updates.

An individual who enters period $t$ of his life with prior $p(=p^t)$ chooses optimal effort vector $e\,(=e^t)$ across occupations to maximize

$$R_t(p) = \max_e \left( b - \delta R_{t+1}(B(p,e)) + H(p,e)\Big(W_t - (b + \delta R_{t+1}(B(p,e)))\Big) - c\left(\sum_i e_i\right) \right). \qquad (4)$$

She can ensure herself the unemployment benefit and the value from continued search. If she finds a job, she looses those but gains the lifetime value of wages ($W_t$). She also has to pay the search costs. For our purposes a two-period model suffices (for which $R_3 = 0$, $W_2 = w$ and $W_1 = w(1+\delta)$).[54] The first period captures the newly unemployed, and the second period the longer-term unemployed.

The alternative interface provides a list of occupations suitable for someone like her. To formalize this, assume that an occupation is only featured on the list if the objective probability $q_i$ of having good job prospects exceeds a threshold $\hat{q}$. In the first period of unemployment this means that for any occupation on the list the individual updates her belief upward to the average of $q_i$ conditional on being larger than $\hat{q}$ (i.e. $p_i^1 = \int_{\hat{q}}^{\overline{q}_i} q_i dQ_i / \int dQ_i$). For occupations that are not on the list, her beliefs decline to the average of $q_i$ conditional of $q_i$ being below $\hat{q}$ (i.e. $p_i^1 = \int_{\underline{q}_i}^{\hat{q}} q_i dQ_i / \int dQ_i$). Obviously these updates also apply if the alternative interface is introduced at a later period of unemployment as long as the individual has not yet actively searched in this occupation.[55] The alternative interface induces an update in belief $p_i^t$ when it is introduced, but given this update problem (4) still applies.

In order to gain some insights into in how an unanticipated introduction of the alternative interface affects the occupational breadth of search, consider for illustration two classes of occupations. Occupations $i \in 1,\ldots,I_1$ are the "core" ones where the jobseeker is more confident and holds first period prior $Q_i = Q_H$ leading to average belief $p_i = p_h$, while she is less confident about the remaining "non-core" occupations to which she assigns prior $Q_j = Q_L$ with average $p_j = p_L$ such that $p_L \leq p_H$. Assume further that core occupations enter the list in the alternative interface for sure (i.e. $\underline{q}_H > \hat{q}$), which means that the alternative interface provides no information content for them. For non-core occupations we assume that there is information content (i.e., $\hat{q} \in (\underline{q}_L, \overline{q}_L)$) so that the alternative interface changes the prior positively if this occupation is featured on the alternative interface and negatively if it is not. For ease of notation, denote by $e_H$ the search effort in the first period in core occupations, and by $e_L$ the same for non-core occupations.

The following results are immediately implied by problem (4): given the search period, the number of core occupations and the current belief about them, there exists a level $\bar{p}$ such that
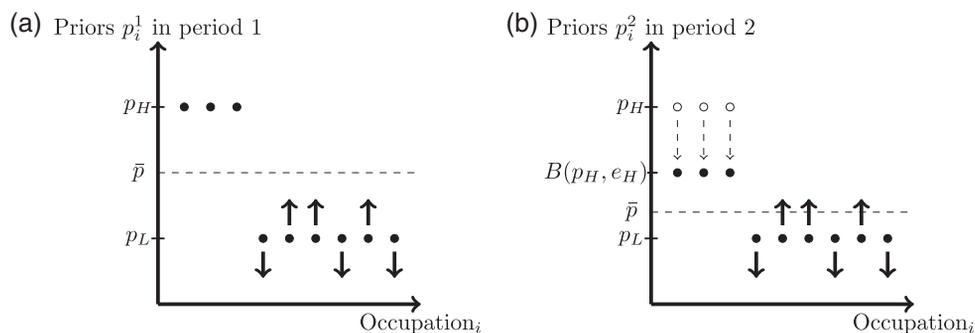
FIGURE 3

Model illustration: narrow search. (a) Narrow search, short-term unemployed: High belief ($p_H$) in first three occupations relative to other occupations ($p_L$). Small changes in beliefs (arrows) do not move beliefs above the threshold $\bar{p}$ to be included into the search. (b) Narrow search, longer-term unemployed: Updating in first three occupations lowers beliefs (dashed arrows). This lowers threshold $\bar{p}$ and additional information moves some beliefs above it and broadens search (solid arrows for occupations 4-10).

the individual puts zero search effort on the non-primary occupations iff $p_i^t \leq \bar{p}$ for each non-core occupation $i$. Intuitively, when the average belief about being a high type in the non-core occupations is sufficiently close to zero, then it is more useful to search in the core occupations and search effort in non-core occupations is zero. More or better core occupations increase the level of $\bar{p}$ since this leads to more search there which drives up the marginal cost of search in non-core occupations.

An individual who is recently unemployed and narrow is depicted in Figure 3(a). She is narrow because her beliefs in her core occupations ($p_H$) are high enough that she does not want to search in the secondary occupations ($\bar{p} > p_L$). This individual concentrates so much effort onto the primary occupations that marginal effort costs are large and exceed the marginal gain from exploring the less likely occupations. In fact, even small changes in the prior $p_L$ induced by the alternative interface—indicated by the thick arrows in the figure - do not move them above the threshold $\bar{p}$.[56] So there is no difference in search behaviour with or without the alternative interface.

In panel (b) we depict our notion of the same individual after a period of unemployment. Her prior at the beginning of the second period is derived by updating from the previous one. After unsuccessful search in the core occupations it has fallen there, as indicated by the lower priors for the first three occupations. Since she did not search in non-core occupations, her prior about them remains unchanged. So beliefs are now closer, and the utility of applying to either of them are also closer and $\bar{p}$ falls.[57] In a model with many rounds beliefs about the core occupations would eventually fall so low that individuals would start searching more broadly even without access to the alternative interface (as we see in our data for the control group, see Section 4). Panel (b) depicts a shorter time frame where $\bar{p}$ falls too little to induce additional search without further information, but enough such that being featured on the alternative interface now moves beliefs about such non-core occupations above it. It thereby induces broader search. Search effort weakly

56. The alternative interface induces small changes if its informativeness is low enough (*e.g.*, $\bar{q}_L - \underline{q}_L < \epsilon$ for sufficiently small $\epsilon$ so that the support of initial beliefs is not very dispersed). We do not explore the alternative case due to its counterfactual implication that already recently unemployed individuals would turn broad with the alternative interface.

57. Additional monetary sanctions for failing to search broader over time would further push down $\bar{p}$ over time.
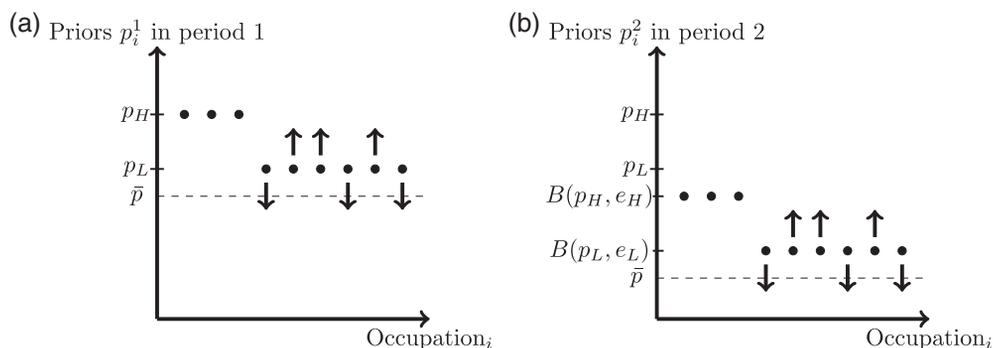
FIGURE 4

Model illustration: broad search. (a) Broad search, short-term unemployed: Beliefs are rather similar and all beliefs are above the search cutoff $\bar{p}$. Changes in beliefs (arrows) move some beliefs below this cutoff, making the person narrower. (b) Broad search, longer-term unemployed: Similar to part (a) but at a lower level of beliefs.

expands, and strictly so if the cost function is smooth, leading to better job prospects.[58] So this rationalizes why longer-unemployed individuals in the treatment group become broader and their number of interviews increases, relative to the control group without the alternative interface. It also implies a weak increase in search effort relative to the control group. At low unemployment durations to the contrary there is little effect.

Figures 4(a) and (b) depict individuals who are already broad in the absence of an information intervention, since the threshold $\bar{p} < p_L$. This could be a recently unemployed individual who started with rather equal priors, as shown in panel (a). Alternatively it could be a person whose beliefs fell over the course of the unemployment spell to a more even level, as shown in (b) (possibly from an initially uneven profile such as in Figure 3(a)). In both cases, beliefs about occupations that are not recommended on the alternative interface might fall so low that the person stops searching there and becomes narrow. Effects on search effort and job prospects are ambiguous: search effort can now be concentrated more effectively on promising occupations which raises job prospects if search effort does not fall too much or even rises; alternatively the negative information on some occupations can lead to such reductions in search effort that overall job prospects fall. Both are theoretically possible.[59] And the offsetting effects can lead to a decrease in breadth without significant employment effects. This can rationalize this empirical finding for initially broad searchers. Thus, the model is able to replicate differential effects by breadth and unemployment duration.

## 7. CONCLUSION

We provided an information intervention in the labour market by redesigning the search interface for unemployed jobseekers. Compared to a "standard" interface where jobseekers themselves have to specify the occupations or keywords they want to look for, the "alternative" interface provides suggestions for occupations based on where other people find jobs and which occupations require similar skills. While the initial costs of setting up such advice might be non-trivial, the

58. Marginal benefits of search in core occupations are not affected by the alternative interface, so optimality implies marginal cost (and thereby search effort) cannot fall. Such effort is only optimal under higher job finding probabilities.

59. With cost functions that induce constant search effort better information must improve job finding. See our previous working paper for the construction of cost functions such that search effort falls and job finding decreases.

intervention shares the concept of a "nudge" in the sense that the marginal cost of providing the intervention to more individuals is essentially costless and individuals are free to opt out and continue with the standard interface. While our intervention has a clear information component that falls within classical economic theory, a major aim of the intervention was to keep things simple for participants so little cognitive effort is required to learn on the alternative interface, which might be considered a nudge element.

We find that the alternative interface significantly increases the overall occupational breadth of job search in terms of listed vacancies. In particular, it makes initially narrow searchers consider a broader set of options, but decreases occupational breadth for initially broad searchers, even though overall the former effect dominates. We find a positive effect on job interviews especially for those which otherwise search narrowly and have an above-median unemployment duration. The effect of unemployment duration is illustrated in our model where those who just got unemployed concentrate their efforts on those occupations where they have most hopes in and are not interested in investing time into new suggestions. If this does not lead to success, their confidence in these occupations declines and they become more open to new ideas.

Some words of caution in line with those in the introduction are warranted. While we find no statistically significant negative effects on job interviews for any subgroup, we cannot rule out that some of them get hurt through less interviews. Moreover, the size of the current study precludes any precise assessment of the effects on job finding, and currently we find no evidence of improvements on this dimension. We have limited information on the types of job found, which jeopardizes our ability to provide a convincing analysis on the duration and quality of jew jobs. At this stage, we can therefore not conclude that the increase in interviews is beneficial. Finally, additional larger-scale roll-out of such assistance would be required to document the full employment effects. The current study does not allow the assessment of equilibrium effects that would arise if everyone obtained information.

With these caveats in mind, our findings suggest that targeted job search assistance to those who otherwise search narrowly and with somewhat longer unemployment duration could be effective, in a cost-efficient way. The programming for the study cost £20,000 ($30,000). If a large-scale website such as Universal Jobmatch would roll out such a scheme for millions of jobseekers, it is obvious that the cost per participant is at the order of a few pence.[60] So any meaningful positive employment effects would swamp the costs. As a first study on job search design on the web, it offers a new route how to improve market outcomes in decentralized environments and hopefully opens the door to more investigations in this area.

---

60. The study also devoted substantial resources (£80,000/$120,000) to attracting participants, compensating participants, and for research assistants.

**Supplementary Data**

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

ALLISON, P. D. and WATERMAN, R. P. (2002), "Fixed Effects Negative Binomial Regression Models", *Sociological Methodology*, **32**, 247–265.

ALTMANN, S., FALK, A., JGER, S. *ET AL.* (2018), "Learning About Job Search: A Field Experiment with Job Seekers in Germany", *Journal of Public Economics*, **164**, 33–49.

ASHENFELTER, O., ASHMORE, D., and DESCHÊNES, O. (2005), "Do Unemployment Insurance Recipients Actively Seek Work? Evidence from Randomized Trials in Four U.S. States", *Journal of Econometrics*, **125**, 53–75.

BEHAGHEL, L., CRÉPON, B., and GURGAND, M. (2014), "Private and Public Provision of Counseling to Job-Seekers: Evidence from a Large Controlled Experiment", *American Economic Journal: Applied Economics*, **6**, 142–174.

BELOT, M., KIRCHER, P., and MULLER, P. (2018), "How Wage Announcements Affect Job Search: A Field Experiment", IZA Discussion Papers 11814, Institute for the Study of Labor (IZA).

BENNMARKER, H., GRÖNQVIST, E., and ÖCKERT, B. (2013), "Effects of Contracting Out Employment Services: Evidence from a Randomized Experiment", *Journal of Public Economics*, **98**, 68–84.

BERGER, M. C., BLACK, D., and SMITH, J. (2000), "Evaluating Profiling as a Means of Allocating Government Services", in Lechner, M. and Pfeiffer, F. (eds), *Econometric Evaluation of Labour Market Policies* (Heidelberg: Physica), 59–84.

BERTRAND, M., DUFLO, E., and MULLAINATHAN, S. (2004), "How Much Should We Trust Differences-in-Differences Estimates?" *The Quarterly Journal of Economics*, **119**, 249–275.

BERTRAND, M. and MULLAINATHAN, S. (2004), "Are Emily and Greg more Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination", *American Economic Review*, **4**, 991–1013.

BLUNDELL, R., DIAS, M. C., MEGHIR, C. *ET AL.* (2004), "Evaluating the Employment Impact of a Mandatory Job Search Program", *Journal of the European Economic Association*, **2**, 569–606.

CARD, D., KLUVE, J., and WEBER, A. (2009), "Active Labor Market Policy Evaluations: A Meta-Analysis" (IZA Discussion Paper Series, No. 4002).

CARD, D., KLUVE, J., and WEBER, A. (2010), "Active Labor Market Policy Evaluations: A Meta-Analysis", *The Economic Journal*, **120**, 452–477.

*Channel 4* (2014), "Why is Government Website Carrying Fake Jobs?" http://www.channel4.com/news/why-is-government-website-carrying-fake-jobs. Posted 07-02-2014. Last accessed 28-09-2015.

*Computer Business Review* (2014), "Universal Jobmatch Here to Stay Despite Fake Job Adverts" http://www.cbronline.com/news/cloud/aas/universal-jobmatch-here-to-stay-but-future-of-provider-monster-is-unclear-4204007. Written by Joe Curtis. Posted 26-03-2014 (last accessed 28 September 2015).

CRÉPON, B., DUFLO, E., GURGAND, M. *ET AL.* (2013), "Do Labor Market Policies have Displacement Effects: Evidence from a Clustered Randomized Experiment", *Quarterly Journal of Economics*, **128**, 531–580.

DINERSTEIN, M., EINAV, L., LEVIN, J. *ET AL.* (2018), "Consumer Price Search and Platform Design in Internet Commerce", *American Economic Review*, **108**, 1820–1859.

FABERMAN, J. R. and KUDLYAK, M. (2014), "The Intensity of Job Search and Search Duration" (Working Paper 14-12, Federal Reserve Bank of Richmond).

GALLAGHER, R., GYANI, A., KIRKMAN, E. *ET AL.* (2015), "Behavioural Insights and the Labour Market: Evidence from a Randomised Controlled Pilot Study and a Large Stepped-Wedge Controlled Trial" (Mimeo).

GAUTIER, P. A., MULLER, P., VAN DER KLAAUW, B. *ET AL.* (2018), "Estimating Equilibrium Effects of Job Search Assistance", *Journal of Labor Economics*, **36**.

GIBBONS, R., KATZ, L. F., LEMIEUX, T. *ET AL.* (2005), "Comparative Advantage, Learning, and Sectoral Wage Determination", *Journal of Labor Economics*, **23**, 681–724.

GIBBONS, R. and WALDMAN, M. (1999), "A Theory of Wage and Promotion Dynamics Inside Firms", *Quarterly Journal of Economics*, **114**, 1321–1358.

GOURIEROUX, C., MONFORT, A., and TROGNON, A. (1984), "Pseudo Maximum Likelihood Methods: Applications to Poisson Models", *Econometrica*, **52**, 701–720.

GROES, F., KIRCHER, P., and MANOVSKII, I. (2015), "The u-shapes of Occupational Mobility", *Review of Economic Studies*, **82**, 659–692.

JOYCE, S. P. (2015), "How to Avoid 5 Major Types of Online Job Scams" http://www.job-hunt.org/onlinejobsearchguide/job-search-scams.shtml. Last accessed 28-09-2015.

KRUEGER, A. B. and MUELLER, A. I. (2016), "A Contribution to the Empirics of Reservation Wages", *American Economic Journal: Economic Policy*, **8**, 142–79.

KRUG, G. and STEPHAN, G. (2013), "Is Contracting-Out Intensified Placement Services more Effective than Provision by the PES? Evidence from a Randomized Field Experiment" (IZA Discussion Paper No. 7403).

KUDLYAK, M., LKHAGVASUREN, D., and SYSUYEV, R. (2014), (Systematic Job Search: New Evidence from Individual Job Application Data) (FRB Richmond Working Paper No. 12-03R).

KUHN, P. and MANSOUR, H. (2014), "Is Internet Job Search Still Ineffective?" *Economic Journal*, **124**, 1213–1233.

LALIVE, R., VAN OURS, J. C., and ZWEIMÜLLER, J. (2005), "The Effect of Benefit Sanctions on the Duration of Unemployment", *Journal of the European Economic Association*, **3**, 1386–1417.

LAUNOV, A. and WAELDE, K. (2013), "Thumbscrews for Agencies or for Individuals? How to Reduce Unemployment" (IZA Discussion Paper No. 7659).

MANNING, A. and PETRONGOLO, B. (2017), "How Local Are Labor Markets? Evidence from a Spatial Job Search Model", *American Economic Review*, **107**, 2877–2907.

MARINESCU, I. (2017), "The General Equilibrium Impacts of Unemployment Insurance: Evidence from a Large Online Job Board", *Journal of Public Economics*, **150**, 14–29.

MARINESCU, I. and RATHELOT, R. (2018), "Mismatch Unemployment and the Geography of Job Search", *American Economic Journal: Macroeconomics*, **10**, 42–70.

MARINESCU, I. and WOLTHOFF, R. (2016), "Opening the Black Box of the Matching Function: The Power of Words" (Working Paper 22508, National Bureau of Economic Research).

MEYER, B. D. (1995), "Lessons from U.S. Unemployment Insurance Experiments", *Journal of Economic Literature*, **33**, 91–131.

MICKLEWRIGHT, J. and NAGY, G. (2010), "The Effect of Monitoring Unemployment Insurance Recipients on Unemployment Duration: Evidence from a Field Experiment", *Labour Economics*, **17**, 180–187.

MILLER, R. A. (1984), "Job Matching and Occupational Choice", *Journal of Political Economy*, **92**, 1086–1120.

MOSCARINI, G. (2001), "Excess Worker Reallocation", *Review of Economic Studies*, **3**, 593–612.

NEAL, D. (1999), "The Complexity of Job Mobility among Young Men", *Journal of Labor Economics*, **17**, 237–261.

ONS (2013), *Internet Access - Households and Individuals, 2013* (UK Office for National Statistics Statistical Bulletin), https://www.ons.gov.uk/peoplepopulationandcommunity/householdcharacteristics/homeinternetandsocialmediausage/bulletins/internetaccesshouseholdsandindividuals/previousReleases.

PAPAGEORGIOU, T. (2014), "Learning Your Comparative Advantages", *Review of Economic Studies*, **8**, 1263–1295.

PATTERSON, C., SAHIN, A., TOPA, G. *ET AL.* (2016), "Working Hard in the Wrong Place: A Mismatch-Based Explanation to the U.K. Productivity Puzzle", *European Economic Review*, **84**, 42–56.

POLLARD, E., BEHLING, F., HILLAGE, J. *ET AL.* (2012), "Jobcentre Plus Employer Satisfaction and Experience Survey 2012" (UK Department of Work and Pensions Research Report No 806).

SAHIN, A., SONG, J., TOPA, G. *ET AL.* (2014), "Mismatch Unemployment", *American Economic Review*, **104**, 3529–3564.

SVARER, M. (2011), "The Effect of Sanctions on Exit from Unemployment: Evidence from Denmark", *Economica*, **78**, 751–778.

*The New York Times* (2009), "Company rarely placed clients in jobs, former employees say" http://www.nytimes.com/2009/08/17/us/17careerbar.html. Written by Michael Luo. Posted 16-08-2009; Last accessed 28-09-2015.

VAN DEN BERG, G. J. and VAN DER KLAAUW, B. (2006), "Counseling and Monitoring of Unemployed Workers: Theory and Evidence from a Controlled Social Experiment", *International Economic Review*, **47**, 895–936.

VAN DER KLAAUW, B. and VAN OURS, J. C. (2013), "Carrot and Stick: How Re-Employment Bonuses and Benefit Sanctions affect Exit Rates from Welfare", *Journal of Applied Econometrics*, **28**, 275–296.

VENN, D. (2012), "Eligibility Criteria for Unemployment Benefits: Quantitative Indicators for OECD and EU Countries" (OECD Social, Employment and Migration Working Papers, No. 131, OECD Publishing).