I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 15452

Do the Long-Term Unemployed Benefit from Automated Occupational Advice during Online Job Search?

Michèle Belot
Philipp Kircher
Paul Muller

JULY 2022

I Z A Institute
of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 15452

Do the Long-Term Unemployed Benefit from Automated Occupational Advice during Online Job Search?

Michèle Belot
Cornell University, University of Edinburgh and IZA

Philipp Kircher
Cornell University, Universite Catholique de Louvain and University of Edinburgh

Paul Muller
Vrije Universiteit Amsterdam, Tinbergen Institute and IZA

JULY 2022

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Do the Long-Term Unemployed Benefit from Automated Occupational Advice during Online Job Search?*

In a randomized field experiment, we provide personalized suggestions about suitable alternative occupations to long-term unemployed job seekers in the UK. The suggestions are automatically generated, integrated in an online job search platform, and fed into actual search queries. Effects on the primary pre-registered outcomes of "finding a stable job" and "reaching a cumulative earnings threshold" are positive, are significant among those who searched at least once, and are more pronounced for those who are longer unemployed. Treated individuals include more occupations in their search and find more jobs in recommended occupations.

JEL Classification: D83, J62, C93

Keywords: online advice, job search, long-term unemployment, occupational mobility, field experiment

Corresponding author:
Paul Muller
VU University Amsterdam
De Boelelaan 1105
1081 HV Amsterdam
The Netherlands
E-mail: p.muller@vu.nl

* We are particularly grateful to Ingeus UK limited for working with us to implement this study. For academic integrity, it was pre-agreed that the authors do not need consent from Ingeus to publish their findings. The authors also gratefully acknowledge funding via UK ESRC grant ES/L009633/1. We thank Andrew Stewart for his dedicated project management and John Allison for building the data collection infrastructure. Authors contributions: All authors contributed to the initial idea, experimental design and writing up of the paper. Kircher and Muller supervised the implementation of the experiment, Muller conducted the data analysis.

1 Introduction

Job search assistance has been shown to increase the chance that unemployed individuals find jobs, but the large personnel costs in administering such interventions often dissipate the benefits (e.g., Behaghel, Crépon, & Gurgand, 2014). Since formal job search has moved mainly online over the last 20 years (Kuhn, 2014) where cheap information transmission is a main advantage, one might want to reconsider this trade-off in this new environment. Belot, Kircher, and Muller (2019) show that automated personalized occupational advice included in a job search engine can significantly raise job interviews at negligible marginal costs. That study was under-powered to study whether the intervention actually raises the probability of employment, nonetheless it showed that effects are concentrated among the longer-term unemployed (above 3 months). The last point opens up interesting new opportunities.

Long-term unemployment is a heavy burden on European economies (e.g., Machin & Manning, 1999; Layard, 2005; Ljungqvist & Sargent, 2008; Duell, Thurau, & Vetter, 2016), and helping to reintegrate such individuals into the labor market is a prime policy target. Interventions on this population tend to be costly, making them particularly suitable for exploring the role of automated advice. Therefore, we adapt the design of Belot et al. (2019) and expand it to provide automated advice during job search specifically to hard-to-place, long-term unemployed individuals.

In cooperation with a private sector company charged with placing long-term job seekers on behalf of the UK government, we integrated our advice into their job search engine. On that engine, individuals can search in two ways: through "keyword search" where they enter a freely chosen search term, or by specifying occupations that they declare as "jobgoals". They are then displayed associated vacancies in their desired geographic area. If they used jobgoal search, for every occupation they chose we suggested up to ten additional occupations. As in Belot et al. (2019), the suggested occupations either experience frequent employment transitions from individuals already working in the initial choice of jobgoal, or they require similar skills. Since some job seekers might lack initial ideas for job goals, we expanded the setup by providing a list of eighteen "no-qualifications" suggestions of occupations with high demand and low skill requirements.

In a randomized controlled trial, 800 participants used the search engine of the

UK partner organization into which we embedded the advice. The treatment group improved on the two pre-registered primary outcomes: the probability of finding a stable job (lasting six months or more) and the probability of reaching a cumulative earnings threshold, both significant at the 10% level. The improvements lie in the 3 - 5 percentage point range, or a 20% to 40% increase relative to the baseline. These effects are comparable, though somewhat smaller, than those on job interviews reported in Belot et al. (2019), which we discuss below. The magnitude is still large given the light touch of the intervention and a target group that was with weak labor market prospects: beyond the long unemployment duration averaging more than five years, one in ten never held a job before, and more than half deal with some form of disability.

The mechanism can be traced directly to the recommendations, rather than to adjustments to the quantity or geographic scope of job search. Individuals in the treatment group do not change the number of searches they perform nor their geographical radius, but when they search via jobgoals they include nearly four times as many occupations (actually, significantly less goals are specified by the individuals, but the additional recommendations increase the total substantially). For the treatment group, this results in a significantly larger share of stable jobs whose occupation is included in our recommended "no-qualifications" occupations.

Effects of this intervention might well have been larger if all individuals had been informed, since only roughly half of them perform a "jobgoal search", a necessary condition to receive suggestions. Job search on the partner platform was not mandatory. Around 600 eligible participants exclusively searched elsewhere and were therefore never informed of their treatment status or exposed to our recommendations. Point estimates and significance are slightly reduced when these individuals are included and results just cease to be significant at conventional levels. Point estimates and significance are strongest among those with above median unemployment duration who search on the platform. This was pre-registered as a heterogeneity analysis. These insights broadly persist when we focus on secondary pre-registered outcome measures comprising "finding any job" (not necessarily a stable job) and "fraction of time employed during first 15 months in the study", as well as in a robustness check of using a duration model rather than our main pre-registered linear probability model. We expect even sharper results in larger samples.

Our study treats only a very small part of the local labor market. As such, there is little scope for spillovers between participants. The scale of this study does not allow us to study broader general equilibrium effects, which have been shown to be relevant in other labor market interventions (Crépon, Duflo, Gurgand, Rathelot, & Zamora, 2013; Gautier, Muller, van der Klaauw, Rosholm, & Svarer, 2018). In particular, it is silent on whether an additional job found within the treatment group comes at the cost of less jobs for others, possibly outside the study population. Given the large differences in job-finding rates between short and long-term unemployed, one would expect large benefits of our intervention even if an additional job for a long-term unemployed comes at the costs of one less job for a short-term unemployed: the generally higher arrival rate of the latter allows faster re-employment and lower cumulative benefit payments compared to prolonged unemployment of the former.

The rest of the paper reviews more of the related literature in Section 2 and provides details about the field experimental setup in Section 3. Section 4 provides the empirical analysis focusing on the primary specification and outcome variables from the pre-analysis plan. It also provides sub-group analysis, studies the underlying mechanisms and considers robustness to secondary outcome measures and empirical specification. Conclusions are offered in Section 5

2 Comparison with the related literature

The paper fits within the broader literature on the effects of counselling (see Card, Kluve, and Weber (2010, 2018) for reviews). To our knowledge, Belot et al. (2019) is the first research study to directly redesign a job search engine and focuses on incorporating automated personalized advice. Their experimental setup comprised roughly 150 people in the control and in the treatment group, each with less than 15% long-term unemployed (above 1 year of unemployment). The vast majority was by design short-term unemployed, with more than three quarters being unemployed for less than half a year. While they find larger effects for those with unemployment duration above 3 months, their participant pool does not allow extrapolation to the long-term unemployed. A particular shortcoming in their setting is that they cannot meaningfully study the effect on job finding, both because they lack access to administrative data and because the sample is too small. While they document large increases in job interviews above 40% on average, and twice the effect size for those with slightly longer unemployment, the fact that these interviews are in occupations that job seekers would not otherwise explore entails the risk that the conversion into actual jobs might be lower.[1] If individuals do not manage to convert interviews in the new occupations, their job finding rate could in principle decrease.

Our study focuses entirely on individuals that are hard to place in the labor market. The vast majority is unemployed beyond one year, and the majority faces some level of disability. Since a non-trivial fraction did not hold a job before, we included the new design element of "no-qualifications" suggestions mentioned in the introduction.

The pre-registered focus of our study is precisely on the final outcome: finding a stable job and reaching enough cumulative earnings. We achieve this via administrative data provided through the partner organization on behalf of the UK Department of Work and Pensions, who acts as final data controller.

The point that active labor market policies are more powerful when directed at the long-term unemployed has already been emphasized in the meta-studies by Card et al. (2010, 2018). They confirm this for training programs that increase work-related skills and for private sector employment subsidies, but find no improvement at any conventional significance level for programs that provide job search assistance (Card et al., 2018, Table 9). Their definition is broad, though, and includes "job search assistance or sanctions for failing to search" (Card et al., 2010, p. 455), which tend to be combined in many interventions. Sanctions have been shown to shift search activity from informal channels to formal channels (e.g., Van den Berg & Van der Klaauw, 2006), thereby reducing overall effects. Belot et al. (2019) document that pure advice unrelated to any sanctions has a positive impact not only directly on the platform where the advice is provided, but also beyond. Arguably this is because knowledge on which occupations to pursue is not only helpful during formal job search on one platform, but also when searching on other platforms or during informal job search. This positive spillover might make pure information interventions more effective. It could be the reason why they find stronger effects at slightly longer unemployment durations where individuals might be discouraged by their existing search strategies and are willing to

[1] Belot et al. (2019) find an overall point estimate of 44% on job finding with significance at the 5% level. For those with above-medium unemployment duration the effect size is 90%, significant at the 5% level.

explore new suggestions. Within our study population we also confirm stronger effects among those with longer unemployment durations.

This seems to be confirmed also in another pure information intervention by Altmann, Falk, Jäger, and Zimmermann (2018) that sent a brochure with generic job search tips to more than 10,000 unemployed in Germany. They find significant positive effects only for those with risk of long-term unemployment. Their study provided the same brochure to all participants with general advice, was offline, and did not link to any particular vacancies. Dhia et al. (2022) report on an even larger study promoting a website that provides broad assistance and advice to job seekers ranging from providing personalized advice on sectors and locations to target; offering step-by-step planning assistance; sending regular reminders and encouragement messages; and providing general tips, such as how to behave during a job interview. The advice provided touches on a much larger range of aspects of job search and the tailored job advice part is not directly linked to any actual vacancies. Effects are very small or absent. Our design provides tailored and individual-specific occupational advice, and makes it cognitively easy by directly linking it to available vacancies that allow job seekers immediately to understand the associated market and internalize the advice through concrete examples. This might explain why our effects are larger and detectable in smaller samples. But these studies share the feature of only providing advice without imposing potential sanctions.

In contrast to studies focusing on more traditional in-person counselling, delivering advice online is substantially less costly. For a recent point of comparison, Cottier, Fluckiger, Kempeneers, and Lalive (2018) evaluate an in-person counselling program targeting long-term unemployed in Geneva, using a comparable sample size. The advice consists of writing job application packages, offering help in finding job vacancies, contacting prospective employers to enquire for vacant positions and referral of candidates. They report effect sizes similar to ours at medium duration, but the monthly cost for each additional participant amounts to 1000 Swiss Francs (roughly equivalent to 1000 Euro) during the first six months and 500 Francs for the rest of the first year.[2] Marginal costs of several thousands of euros for advice to an additional job seeker have

[2]The impact on employment is around 5 percentage points during the first year, similar to our finding. It reverses to zero in the second year. We do not see a reversal in our pictures, even if we extend them to 15 months, but we loose observations because not all unemployed are observed for this duration and we therefore do not explore this further in this paper.

also been reported in Behaghel et al. (2014) and other studies cited therein, and are the reason that the cost-benefit analysis for in-person advice through private sector agencies in their setting does not come out positive. The marginal cost in our online job search setting is essentially zero. Despite several months of designing the advice and their integration into the job search site, even the average cost per job seeker is trivial compared to the costs of in-person studies.

3 Experimental details

3.1 Context and Partner organization

The experiment was implemented in the United Kingdom at the start of 2019, in collaboration with a large, international provider of support programs for unemployed job seekers towards finding employment. In the UK, job seekers initially sign up with a local public job center that offers a range of services, including regular meetings with a caseworker, whose role is to both monitor the job seekers' search efforts and provide advice.

After a year registered as unemployed, a job seeker has to transfer to one of the government's workfare programmes that are run by private companies.[3] In 2017, England and Wales rolled out a specific program called "The Work and Health Programme" (WHP) that predominantly helps disabled people, as well as the long-term unemployed, and certain priority groups (known as early access groups) to enter into and stay in work. The WHP offers tailored support to meet individual needs. Examples of the type of support available may include participants having personalised support with regular face to face contact, liaison with local health and well-being professionals and with dedicated employer experts with knowledge of the local labour market and job opportunities. Providers will support participants for up to 15 months. This may be extended for a further 6 months to provide in-work support, to a maximum total of 21 months on the WHP.

The providers are paid a service delivery fee as well as outcome related payments when a person reaches a specified level of earnings once in employment, or records

[3]Note that outsourcing of counselling and placement services is becoming increasingly common in a number of countries like the UK, Germany or France. Whether this results in more or less effective support to job seekers is debated (Bennmarker, Grönqvist, and Öckert (2009), Behaghel et al. (2014), Krug and Stephan (2013))

6 months of being in self-employment. The private company we collaborated with is one of the largest private providers of support services for job seekers. It was initially founded in Australia in the 1980s and has since then expanded to a number of countries. The company specializes in programs for job seekers that require special assistance, such as long-term unemployed and also offers rehabilitation services to injured workers, supporting their return to work.

Note that in terms of financial assistance, UK job seekers are eligible to receive a "Job Search Allowance" (JSA) for a limited period of time (up to six months) provided they meet the eligibility criteria.[4] Once the JSA stops, a job seeker can apply for Universal Credit. The amount of support they receive is based on circumstances, such as health condition or disability, income and housing costs.

At the time of the experiment, the UK unemployment rate was at a historical low of 3.9%. The region North-West of England, where the experiment was conducted, had a level of unemployment comparable to the national level (3.6%)[5]. The share of long term unemployed (i.e. searching for a job for longer than a year) in 2019 was 28%.[6]

3.2 Occupational suggestions

The online intervention consisted of integrating occupational suggestions into the job search interface. This type of intervention fits with developments in "online recommendation systems" that have recently burgeoned across various markets (Bergemann & Ottaviani, 2021). To our knowledge and aside from our previous study (Belot et al., 2019), the only other study that evaluates an algorithmic recommendation system focuses on a specific online gig platform (Horton, 2017).

The bulk of current online recommendation systems rely on machine learning algorithms, which are based on data on users' behavior on the platform. The main challenge job search platforms face is that they typically do not observe the end outcome of the

[4]At the time of the experiment, the criteria included being at least 18 years old (some exceptions apply for 16 and 17-year-olds), being available to work, being below the State Pension age, not in full-time education, be in England, Scotland, or Wales, not working at the moment (or less than 16 hours per week as an average), be single (or having a partner who is working less than 24 hours a week as an average), not having an illness or disability that stops you from working and having no more than £16,000 in savings (including any savings that a partner has).

[5]source: UK Office Of National Statistics

[6]source: OECD statistics

search, i.e. whether people found a job or not.[7] We therefore opted for alternative methodologies to develop the advice algorithms and ensure that the advice we provide relates to potentially promising occupational transitions.

We rely on two methodologies to generate relevant suggestions. A suggestion is a combination of a "base occupation" and a "suggested occupation". The base occupation is typically the occupation that a job seeker is searching for, and it might be the occupation corresponding to their previous job or to their qualifications. The suggested occupation is the occupation that is suggested as an alternative for the job seeker. The first methodology creates suggestions from common occupational transitions (as observed in labour market surveys or obtainable retrospectively from a CV database). The second methodology creates suggestions that require (largely) the same set of skills as the base occupation. Each methodology uses the UK SOC 2010 (4-digit) occupational classification.

The first methodology is based on common occupational transitions in the UK Labour Force Survey and the UK Household Longitudinal Study ("Understanding Society"). Both surveys capture employment histories and we identify occupational transitions by changes in occupation code occurring with a job transition. We constructed a transition matrix, which counts how often each potential occupational transition occurred in the data. For each base occupation, the 10 most common transitions were selected to serve as the first set of suggestions. Note that due to the limited size of the underlying surveys, for many occupations there were less than 10 suggestions available.[8]

Our second methodology is based on transferable skills. Occupations that require similar skills are likely to offer feasible transitions. We used the website O*NET, which provides for each base occupation around 10 alternative occupations that require similar skills.[9] For each occupation, all related occupations were ranked based on the

[7]Naya et al. (2021) discuss in depth other problems of recommender systems, in particular how to take into account congestion: for most jobs only one job seeker can fill it, so recommending the same job to many people can be problematic. They provide algorithms to relieve this and to incorporate preferences, and circumvent the problem of knowing the outcome of search by using French administrative data.

[8]As an indication, out of the 369 base occupations, there was at least 1 suggestion for 285 occupations and at least 3 suggestions for 218 occupations.

[9]O*NET information uses the US occupational classification, which we transfer into the UK 2010 SOC using a crosswalk provided by the UK Department of Education. The two classifications do not map one to one, thus there are typically several US occupations linked to one particular UK occupation. In those cases we use the US employment shares for each occupation to select the most populous occupations.

UK employment shares and the 10 largest were kept as suggestions. As in the first methodology, this did not provide 10 suggestions for each occupation, but there were at least 2 suggestions for 341 occupations. The final set of suggestions combines the suggestions from the two methodologies.[10]

As a second element of the intervention, we also offered a number of 'no-qualifications' occupations for those job seekers with little working experience. We first shortlisted occupations that the US Bureau of Labor Statistics associate with no formal education requirements. We then selected the 18 occupations with the largest number of vacancies in the UK public vacancy database (Universal Jobmatch) in March 2016. For each of these 18 occupations, a short description of the job and main tasks was provided, as well as an illustrative picture (see the Online Appendix for an example).

3.3 Online implementation

Our aim was to make the occupational suggestions easily accessible and easy to incorporate into job search. To do so, we worked together with the developers of the job search platform that was used by the job seekers participating in the program. On the platform job seekers can search for vacancies, apply for jobs, and keep track of their progress by tracking jobs and applications. Users of the platform first create an account, which is typically done during a meeting with their case worker. Job seekers are encouraged to use the platform, although there are no strict obligations.

Once registered, users can search for jobs by creating 'Job Goals', which are the occupations in which they aspire to find work. After selecting their job goals, they can perform a search query over these job goals, and additionally add other restrictions based on geographical location. Job goals are part of the user's account and are automatically saved for future job search. They can be removed or changed at any time. The suggestion treatment is implemented at the stage of selecting job goals. When a job seeker selects an occupation as jobgoal, the platform shows the list of suggestions of alternative occupations, each of which can be ticked or unticked depending on preferences. When performing the next job search query, the selected suggestions are automatically incorporated, although the search page also provides a simple button to

[10]The final set of suggestions contains all suggestions that appear in both methodologies. Next we extend the list to a maximum of 10 suggestions using the highest ranked occupations from either methodology.

exclude all suggestions. Job seekers can also search by simply entering keywords, in which case occupational suggestions are not provided. A screenshot of the job search page is provided in the Online Appendix.

After performing a job search query, matching vacancies are displayed as is commonly the case on online platforms: job titles and locations are shown for a limited number of jobs per page. The user can click to proceed to the next page, and can also click on individual ads to view all information about the job. Search results can be ordered by distance, publication date, or 'relevance' (based on how closely the job matches the search keyword).

3.4 Randomization, Timeline, Data Collection, and Power

We test our intervention in a randomized experiment. The participants were all job seekers that participated in the 'Work and Health Programme' (WHP) offered by the support provider within the greater Manchester area in England. As explained above, this support program was targeted specifically at job seekers that have a health condition or disability, or had been unemployed for a long time (81% of our sample was unemployed for more than one year). As part of the program, participants were encouraged to register on the job search platform, and use it for their daily job search and keep track of their progress. The date of registration on the platform determines selection into the randomized experiment, with all job seekers registering after January 1st 2019 being selected. From early 2020 onward the platform was phased out and replaced by a new platform, resulting in essentially no new participants after March 2020. We present the density of registration dates in Figure A.1 (panel a) in the Appendix, showing that the bulk of our participants registers in the first half of 2019. Also the number of search activities on the platform became negligible towards March 2020, implying that our study essentially ended before the Covid pandemic started.

The start of the experiment was delayed by several months relative to the pre-analysis plan. Job search is only a small part of the overall platform that customers can use to interact with the partner organization, and overall implementation issues led to a delayed roll-out, as well as to the premature phase-out of the overall platform. These issues also implied that the platform was not rolled out in other geographical areas of the WHP. The shorter duration and smaller scope led to a smaller sample size than initially expected.[11] If the effect size on job finding in our intervention is similar to the 44% point estimate in Belot et al. (2019) for additional job interviews, keeping all other parameters similar to those from our pre-analysis plan, we would need 500 participants to be sufficiently powered.[12] At half that effect size we would need 1500 individuals to detect this effect with 10% significance at the usual 80% power. Both are within the sample size of this study. We are under-powered to detect smaller effect sizes.

At the time of registration, users were presented a consent statement explaining that an academic study was taking place to improve the job search process (see Online Appendix for the consent statement). At this stage they were asked to consent to (anonymous) data being collected and processed. They were informed that consent was voluntary and not required for using the platform. If they refused, they would proceed to the platform and would not be of the experimental sample. If they provided consent, the platform assigned them randomly to either the treatment group (60%) in which case they would see our suggestions, or the control group (40%) in which case they would use the platform without suggestions.[13]

For participants in the experiment, the platform collected job search statistics which were transferred into a research database using an API. In particular, we obtained data on all jobgoals and job search queries (both jobgoal-searches and keyword searches) including search parameters and the returned set of vacancies. These data were, after completion of the study, linked with data from the provider's CRM on individual characteristics and information on job finding and job interviews. This linking was performed on January 15 2021, implying that job finding information is right censored at this date.

[11] In the Online Appendix we provide the pre-analysis plan and an overview of the deviations from it in terms of data collection and analysis.

[12] In our pre-analysis plan we computed power under the assumption of a base-rate of job finding of 28 percentage points in the control group, and type-I error of 5% and a type-II error of 20% (i.e., 80% power). Belot et al. (2019) reported a point estimate for the increase in the treatment group of 44% for job interviews. Using this number as the effect size for job finding, and accounting for an enrollment ratio of 1.5 (60% of our participants get treated), leads to the reported number. For robustness, all else equal, using the realized job finding rate of our control group (16% after 10 months) as the base would require roughly twice the sample size. On the other hand, all else equal, using the 90% effect size that applies to the longer-term unemployed (above three months) in Belot et al. (2019) would cut the required sample size by roughly three quarters.

[13] The 60-40 split was agreed upon due to the provider's preference to offer the enhanced job search platform to as many job seekers as possible.

Table 1: Balance table

	All job seekers			Searched on platform (57%)		
	Control	Treatment	P-value	Control	Treatment	P-value
Female	0.41	0.39	0.65	0.40	0.41	0.91
	(0.49)	(0.49)		(0.49)	(0.49)	
Non-white ethnicity	0.25	0.26	0.76	0.26	0.26	0.89
	(0.43)	(0.44)		(0.44)	(0.44)	
Married/Cohabiting	0.10	0.10	0.88	0.11	0.11	0.89
	(0.30)	(0.30)		(0.31)	(0.31)	
Children	0.15	0.17	0.37	0.16	0.20	0.20
	(0.36)	(0.37)		(0.37)	(0.40)	
Jobseeker's allowance	0.40	0.38	0.48	0.41	0.39	0.58
	(0.49)	(0.49)		(0.49)	(0.49)	
Unempl. dur. (months)[a]	62.20	62.36	0.96	64.84	61.44	0.43
	(57.76)	(58.28)		(58.87)	(58.10)	
Never worked before	0.10	0.08	0.32	0.10	0.07	0.24
	(0.30)	(0.27)		(0.30)	(0.26)	
Disability	0.54	0.56	0.46	0.57	0.56	0.77
	(0.50)	(0.50)		(0.50)	(0.50)	
N	525	869		298	502	

***, **, and * indicate significance at the 1, 5, and 10 percent critical level. P-values refer to a test for equality between control and treatment group. [a] Unemployment duration in months at the time of starting on the WHP programme, excluding individuals that never worked before.

4 Empirical analysis

The randomized assignment allows for a straightforward comparison of the control and treatment groups. We start by presenting a balance table with descriptive statistics on our sample. Next, we show differences in job search success measures, as defined in our pre-analysis plan. Afterwards we briefly show evidence on job search strategies.

4.1 Descriptive statistics and balance

Table 1 shows pre-intervention balance between control and treatment groups in columns 1 and 2. The groups are perfectly balanced with no significant differences. Note that there are more treated (869) than control group (525) participants, as the experiment implemented a 60% treatment assignment. The sample is fairly disadvantaged relative to regular unemployed job seekers: more than half have some disability, and the average unemployment duration (at the time of starting in the program) is more than 1 year (and this excludes a 10% share that reports to have never worked before).

Out of all participants who registered on the platform, 57% used it to perform at least one search query. The randomization into treatment and control occurred at the *time of registration*, but the differences on the platform due to treatment assignment are only visible to users when they select job goals and perform a search query. Therefore we can safely restrict our sample to those job seekers that searched at least once on the platform. This reduces noise, as our intervention can only affect those participants that actually received the advice during a search query. Columns 4 and 5 of Table 1 confirm that the subset of participants that searched on the platform are very similar to the full sample. In addition we do not find any statistically significant differences between treated and control group within those that searched on the platform.

4.2 Primary Outcome: Job finding

The goal of the intervention is to provide suggestions of alternative occupations, which should increase the chances job seekers find employment. We test our hypothesis by comparing job finding outcomes between control and treatment groups. Our pre-analysis plan specified the empirical methodology and two primary outcomes: the duration until finding a stable job (defined as a job that lasts at least six months) and the duration until reaching an earnings threshold. These outcome measures are complementary. The first outcome ensures that one only counts jobs that provide some medium-term perspective for the job seeker. While being the first-order measure of interest, the definition of a stable job does not consider the hours or earnings and neither does it include cases where an individual is employed for a continuous time period but switches jobs in that period. Therefore we also use a second primary outcome which is the duration until having reached a (cumulative) earnings threshold of 3,415 pounds (which equals working 6 months for 16 hours per week at the minimum wage). This earnings threshold corresponds to a performance measure used to evaluate the service provider.

Our empirical model is simple. For t months after registration on the platform, define Y_i^t to be an indicator equal to 1 if the duration until finding a job (or reaching the earnings threshold) is smaller than t for individual i and 0 otherwise. We then estimate

$$Y_i^t = \alpha + \beta T_i + \gamma X_i + \epsilon_i, \tag{1}$$

with T_i a treatment group indicator and X_i a set of time-invariant control variables that includes gender, ethnicity, disability, marital status, children, job seeker's allowance receipt, education, unemployment duration at the time of starting in the program, and quarter of registration dummies.[14] We estimate equation (1) for $t \in 1 - 12$, to examine how the difference in job finding between control and treatment groups changes over the first 12 months.

Finding a stable job

For each individual we observe all jobs found within our observation period including the starting date and end date. We assume the duration is right-censored for jobs without end date.[15] We exclude all jobs that started prior to registration on the platform, as these cannot be related to our treatment. We present regression estimates visually in Figure 1 (with 90% confidence intervals), where t runs from one month after registration to 12 months after registration. As mentioned, we include only individuals that searched at least once on the platform. The y-axis depicts the difference in the share of individuals that started a stable job between the control and treatment groups, after controlling for the above mentioned covariates (the coefficient β in equation 1). In panel (a) we find that the share of stable job finders is larger in the treatment group and grows over time to about 5%-points. The difference is statistically significant at the 10% level after 1-2 months and at longer durations (8-10 months). The magnitude of the estimates (around 3-5 percentage-points) can be compared against the control group stable job finding share after 10 months, which is 16%, implying a 19-31% increase.

When including all participants (also those that never searched on the platform), we find a similar positive difference between treatment and control group (panel (a) in Figure 2). The difference is never statistically significant, which is likely due to the relatively large fraction (40 %) of job seekers who never use the platform.

[14]The pre-analysis plan specified that additional fixed effects would be included for the office where the individual registered. Unfortunately we did not receive these data.

[15]It is possible that the end date is missing for some jobs, which would imply that we overestimate the stability of some jobs. But since there is no reason to expect that this differs by treatment status, this should not matter for our results.

Figure 1: Coefficients linear regressions: participants that searched on the platform

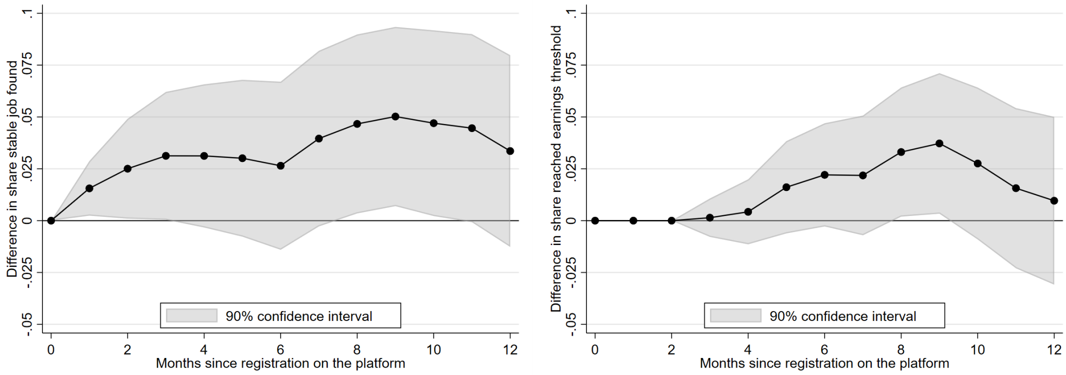

(a) Finding a stable job

(b) Reaching earnings threshold

Figure 2: Coefficients linear regressions: all participants

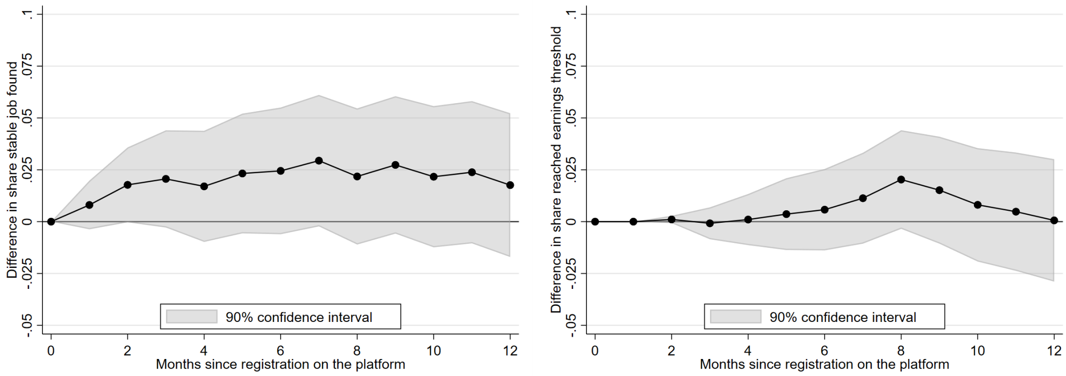

(a) Finding a stable job

(b) Reaching earnings threshold

Reaching an earnings threshold

The duration until reaching the earnings threshold is based on data on cumulative earnings, possibly across jobs, per individual. The data is collected by the provider and is part of their performance evaluation. Earnings are calculated using the reported number of weekly working hours and the hourly wage. As above, we assume jobs with no observed end date to last at least until the end of our observation period. In panel (b) of Figure 1 we show the difference in the share of individuals that reach the earnings threshold at varying durations since registration. The number of individuals that reach the threshold within the first couple of months is negligible and correspondingly the difference between treatment and control only opens up from five months onwards. The difference grows to a 3%-points higher share in the treatment group after 9 months (significant at the 10% level). As a comparison, 8% of the control group reached the earnings threshold after 9 months, implying a 40% increase. In panel (b) of Figure 2 we include all participants and find a qualitatively similar picture, although the difference is slightly smaller and not significant.

Heterogeneity analysis: unemployment duration

Our pre-specified hypotheses state that the impact of occupational advice is expected to vary by initial unemployment duration and initial 'breadth of search' (motivated by the findings in (Belot et al., 2019)).[16] Unfortunately our data does not contain information on job search prior to the intervention, making it impossible to differentiate between those who searched broadly and those who searched narrowly prior to our study. We are able to split the sample by unemployment duration as measured at the time of starting in the program. The hypothesis that unemployment duration matters follows from the idea that with longer duration job seekers may become more willing to explore alternative options, which can rationalize the increased impact on this group e.g. in Belot et al. (2019). We split the sample at the median duration, which is 18 months of unemployment.[17] The results are presented in Figure 3 for those with long unemployment duration. The difference in finding a stable job (panel a) is larger for this group, with those in the treatment group having a 7% points higher job finding after 9 months. The difference is statistically significant at the 10% level for all durations up to 12 months. In panel (b) we find a similar difference for reaching the earnings threshold: the treated group reaches the threshold at a higher rate, with a significant difference from four months after registration onward. The flip side is that those with shorter unemployment duration do not appear to have benefited from the advice (Figure 4). The control and treated groups do equally well in securing stable jobs (panel a) and reaching the earnings threshold (panel b).

Larger effects of interventions for those with longer unemployment duration have often been observed in the literature on active labor market policies, as discussed in the

[16]Note that the pre-analysis plan also indicated to separate the analysis by whether individuals had a disability if treatment effects differed substantially across this dimension. Since this is not case, we abstain from exploring such heterogeneity further.

[17]Data on unemployment duration is binned, such that we have 41% with unemployment duration of at most 18 months, and 59% with unemployment duration of more than 18 months. Those who have never worked before are included in the long duration sample.

Figure 3: Coefficients linear regressions: participants that searched on the platform and have long unemployment duration

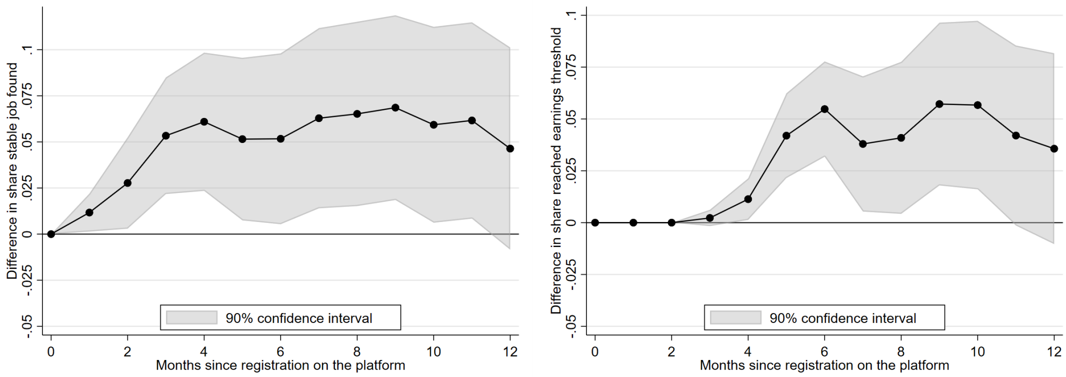

(a) Finding a stable job (b) Reaching earnings threshold

literature review. In our case, those with shorter unemployment durations are likely to have moved more recently from the standard system of the UK department of work and pensions to specialized services such as those offered by our partner organization. Such a transition usually happens around 12 months, but the timing depends on many factors including other obstacles to placement such as disabilities. Such individuals might focus more on other changes and pay less attention to our advice. Those with longer unemployment durations are also those who tend to be in such special programs (though not necessarily with our partner organization) for longer. Our advice might be the main novelty for such individuals, and might receive more notice. Unfortunately we do not have data to investigate this more deeply, but it is noteworthy that those with long unemployment durations do not seem to be the ones resistant to change.

4.3 Secondary Outcomes

Occupational Mobility

Evidently, the key mechanism for why the advice increased job finding is the breadth of occupations job seekers considered. We next examine how the advice affected occupational transitions. For the participants that found jobs, we observe the job title. Using an automated classification tool (CASCOT), we assign Standard Occupational Classification Codes (SOC's) to each job. We can then compare the nature of transitions between treated and control groups, conditional on job finding. Given the higher rate

Figure 4: Coefficients linear regressions: participants that searched on the platform and have short unemployment duration

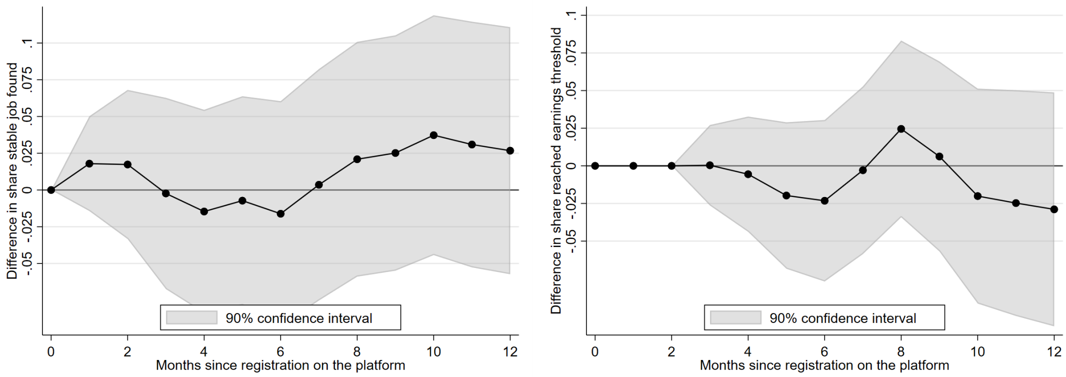

(a) Finding a stable job

(b) Reaching earnings threshold

Table 2: Is the occupation of the first stable job one of the 'no-qualifications' suggestions?

	Searched on platform		All participants	
	(1)	(2)	(3)	(4)
Treatment group	0.19**	0.16*	0.12*	0.12*
	(0.084)	(0.091)	(0.064)	(0.065)
Observations	151	151	256	256
Mean outcome	0.52	0.52	0.52	0.52
Control variables	No	Yes	No	Yes

Robust standard errors in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level.

of job finding in the treatment group that we documented, there may be a selection problem that we return to below. Figure A.4 in the Appendix shows the distribution of jobs across 'Major Groups' (first-digit occupations) by treatment status. The distributions look fairly similar, with the majority of jobs in 'Elementary', 'Sales and Customer Service' and 'Personal Service'. The low share of jobs within higher-skilled categories reflects the specific population taking part in the assistance program that we consider.

To study whether the occupational suggestions affect the types of jobs that were found, one would ideally compare the occupation of the new job with that of the previous job. Unfortunately this is not feasible, because we do not have information on the previous job. Instead, we assess whether the new jobs are more likely to be among the occupations that we suggested as 'no-qualifications' suggestions (see section 3.2). These no-qualification suggestions were directly added as jobgoals by 23% of the treatment group (among those that searched at least once on the platform). Of course, more users may have seen them and thereby internalized the information.

First, we find that a large share of those that find stable work, do so in one of our suggested no-qualifications occupations (52%). This reflects that the population of job seekers is relatively low-skilled and that the no-qualifications suggestions are occupations that make up relatively large shares of the working population. In Table 2 we regress an indicator for the first stable job being in one of the suggested no-qualifications occupations on a treatment group dummy. Column (1) shows that among those that searched at least once, the treatment group is much more likely to have found a job in such an occupation than the control group (19 %-points higher probability, significant at the 5% level). Adding the set of control variables reduces the coefficient only slightly (column (2)), but leads to a slightly larger standard error. Including all individuals even if they did not search on the platform still displays a large impact (12 %-points higher probability, and significant at the 5% level; columns (3) and (4)). Overall, we find that exposure to the suggestions leads to a substantial increase in the likelihood that the first stable job is in one of the suggested no-qualifications occupations.

One may worry that the increased likelihood of finding employment in one of these suggested occupations may have detrimental consequences if they represent jobs with, for example, lower wages. The fact that we also find an increased likelihood of reaching the earnings threshold refutes such worry as it shows that the increased job finding in the treatment group does not reduce their earnings. Finally, we show the composition of new jobs between suggested and other occupations for the treatment and control groups in Figure 5 (for those that searched on the platform). Indeed the increased share that finds a job in our suggested occupations leads to a slight reduction in the share that finds a job in any other occupation (similar for the full sample in Figure A.3 in the Appendix).

Other measures of job finding success

We pre-specified two other measures of job finding success as secondary outcomes of interest: (1) the duration until starting the first job and (2) the fraction of time employed during the first 15 months. Clearly these are correlated with our primary outcome measures, although they may capture slightly different dimensions: starting a job that does

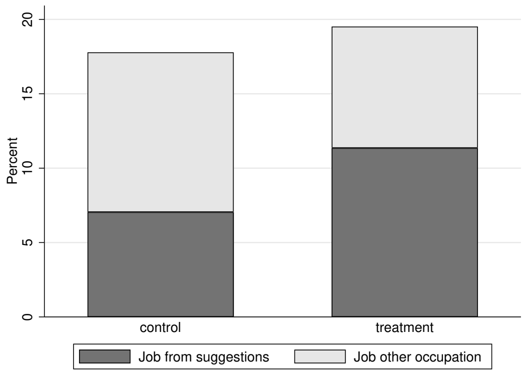

Figure 5: Occupations of the first stable job found: those that searched on the platform

not last for six months and neither pays enough to reach the earnings threshold would not be included as a success in the previous findings. In Appendix Figure A.2 we present the results from estimating equation (1) for the duration until starting the first job. We exclude jobs that last less than three days. Results are similar to those found for the primary outcomes: the job finding rate is higher for the treatment group overall and in particular when focusing on those that searched at least once, although confidence intervals are wider which impacts statistical significance (panels a and b). Among those that used the platform, the long-term unemployed do benefit significantly from the occupational suggestions (panel c), while again there is no difference for the group with relatively shorter unemployment duration (panel d).

The fraction of time employed during the first 15 months is a measure that takes the duration of jobs into account, and aggregates information across multiple jobs that an individual might hold. For those that enter the study less than 15 months prior to the end of our observation period, we compute their fractions within the remaining observed time period. We regress the fraction employed on a treatment dummy and the same set of control variables. Results are presented in Table A.3 in the Appendix. The treatment group experiences a longer fraction of time employed, with the effect larger among those that searched on the platform, and even larger among those with relatively long unemployment durations. The difference is only significant for the latter (at the 10% level).

We conclude that these two secondary outcomes of interest reinforce our primary findings. Job seekers in the treatment group appear to do better in terms of employment, with the impact larger for those that used the platform, and largest for those with long unemployment durations.

4.4 Robustness: duration model

We evaluate the impact of our intervention using an alternative model to assess robustness of the main findings. We estimate a duration model (Cox proportional hazard model). This specification has the advantage of using more detailed information on job finding: the exact duration is used as the outcome rather than an indicator variable and right censoring is accounted for. We include the same set of control variables as in the baseline analysis in all models.

We examine durations related to our two primary outcomes of interest: duration until finding a stable job and duration until reaching the earnings threshold. Results for the group of individuals that used the platform at least once are presented in Table A.4, for the full sample in Table A.5, and for those with relatively short and long initial unemployment durations in Tables A.6 and A.7. The Cox proportional hazard model estimates a constant proportional effect of the treatment on the hazard rate. To investigate differential effects of the intervention at different durations, we censor the outcome at different durations. The results are in line with our baseline findings: the treatment increases the rate at which participants find stable jobs and the rate at which they reach the earnings threshold. For the full sample the difference is not significant, but increased job finding in the initial months is statistically significant for those with longer initial unemployment that searched at least once.

4.5 Job search activity

Since we have data on searches on the platform, we can examine in more detail the impact of the treatment on job search behavior. Because usage of the platform was limited with a mean of less than three search queries per individual, we provide simple mean comparisons, rather than estimating panel models that consider the evolution of search over time.[18] Table 3 compares job search behavior between control and treatment groups for those that searched at least once on the platform. The top panel shows statistics on the jobgoals that users saved in their profile. First, we find that in the treatment group 52% saved at least one jobgoal in their profile, which is a proxy

[18]Our pre-analysis plan specified an individual fixed effects model that analyses how job search activities develop over time. This is not feasible with the number of observations that we have.

Table 3: Job search behaviour: only those that searched at least once on the platform

Variable	N	(1) control Mean/SD	N	(2) treatment Mean/SD	T-test P-value (1)-(2)
Jobgoals:					
Saved at least one jobgoal	298	0.41	502	0.52	0.00***
		(0.49)		(0.50)	
Mean nr saved jobgoals	298	1.40	502	1.29	0.50
		(2.63)		(1.93)	
Saved at least one manual jobgoal	298	0.41	502	0.36	0.21
		(0.49)		(0.48)	
Mean nr manually saved jobgoals	298	1.40	502	0.91	0.00***
		(2.63)		(1.71)	
Saved at least one no-qual. jobgoal	298	0.00	502	0.23	0.00***
		(0.00)		(0.42)	
Mean nr saved no-qual. jobgoals	298	0.00	502	0.37	0.00***
		(0.00)		(0.86)	
Search queries[a]:					
Total nr of jobgoal searches	298	1.63	502	1.57	0.91
		(7.56)		(7.80)	
Total nr of keyword searches	298	3.09	502	2.98	0.83
		(6.32)		(7.35)	
At least one suggestions-search[b]	298	0.00	502	0.44	0.00***
		(0.00)		(0.50)	
Average nr of included jobgoals (excl sug's)[c]	113	2.97	222	2.17	0.00***
		(2.30)		(1.54)	
Average nr of included suggestions[d]	113	0.00	223	7.79	0.00***
		(0.00)		(5.21)	
Average radius in jobgoal searches	113	8.94	223	13.93	0.43
		(4.11)		(66.90)	
Average radius in keyword searches	234	8.50	382	8.41	0.76
		(3.96)		(3.45)	
Search query results:					
Median nr of results in jobgoal searches[e]	113	1460.65	223	5230.49	0.00***
		(3324.77)		(12997.96)	
Median nr of results in keyword searches[e]	234	9414.19	382	12834.30	0.24
		(28783.16)		(38252.30)	

The sample sizes N refer to the number of individuals that are included in the respective means. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level. [a] All variables are first averaged by individual (across search queries) and then averaged across individuals. [b] Indicator for whether the individual performed at least one jobgoal search with the suggestions included. [c] The average number of jobgoals that were included in the jobgoal search queries (this does not include the suggestions that may have also been included in the search query). [d] The average number of included suggestions within jobgoal searches. [e] The average over the median number of search results per individual.

for having seen the suggestions that we offered. Given that the share is significantly lower in the control group (41%), our suggestion appear to have stimulated participants to use the jobgoal functionality. Across all individuals, the treatment group saved on average 1.29 jobgoals, of which 0.91 are 'manually added' jobgoals and the remaining 0.37 jobgoals come from the no-qualifications suggestions. In the treatment group 23% saved at least one of these 'no-qualifications' suggestions as a jobgoal. The control group saved 1.40 jobgoals on average, which are all 'manually added'.

The second panel displays statistics about job search queries performed.[19] There is no difference in terms of the number of searches performed (both for keyword searches and jobgoal searches), but for the treatment group, 44% included some of the suggestions at least once in a search query. When considering the jobgoal searches, we see that the treatment group included fewer jobgoals (2.17 vs 2.97), which might be a response to the inclusion of additional suggestions generated by our treatment (7.79 on average). As a result the total number of included occupations in their searches is much higher. Geographically we find no significant differences, with the average radius around 10 miles. The effect of the included suggestions in jobgoal searches does translate into significantly more search results, while for keyword searches there is no difference between the groups (as expected).

Overall, our treatment induces participants to increase the number of occupations in their search queries. They achieve this with less jobgoals chosen themselves by including a large number of suggestions. Neither the number of job goals or keyword searches nor the average geographical radius is affected. The equivalent statistics for the full sample (including those that never searched on the platform) are similar and presented in Table A.9 in the Appendix.

5 Conclusion

Job search assistance has long been a fruitful instrument of active labor market policies. Providing advice is particularly appealing because it uses no coercion: individuals are free to adopt this advice if they consider it useful. Unfortunately traditional methods of advising job seekers are very labor intensive.

Integrating personalized advice online right at the point where it is useful has several advantages. It is relatively easy to understand and implement, and the results can be immediately seen by the target audience. This makes it cognitively easy to absorb the

[19] Each variable is first averaged across queries by individual, before averaging across individuals.

advice, and bears the potential that the advice translates into action. Once the advice is absorbed, it also allows the audience to use that knowledge in other circumstances, for example in talking to friends and family.

Here we explore this concretely in the context of occupational advice to long-term unemployed job seekers. A priori it might seem daunting to use simple algorithmic advice on a population that is out of work for more than a year, where more than half face issues of disability, and one in ten has never worked before. Nevertheless, this advice has measurable effects on their chances of successful reintegration into the labor market. Relative to earlier literature that only documented increased job interviews (and did not focus on the long-term unemployed), it is crucial to examine success in actual job placement, as interviews in novel occupations might not convert easily into jobs. Here we do not have information on job interviews, but can directly trace the final labor market consequences with administrative data.

Overall, this paints a promising picture of using targeted advice during online search. We see three obvious issues that future research might address. First, the interaction between human advice and automated advice. In our setting all participants also receive the standard package from the partner organization, with advisors unaware of treatment status. This package includes substantial resources to remove barriers to re-employment (motivational and emotional support, help in solving disability related hurdles, mobility assistance for those far from employment centers, training to get a drivers licence, etc) which might well complement online activities. Second, online advice can be rolled out at large scale rather easily, but requires access to a large target population. In our case we initially expected ten times the sample size, but were limited by the implementation of the partner organization.[20] Larger roll-outs allow for tighter confidence intervals, and can also uncover general equilibrium effects if there is differential role-out across regions. By treating many individuals in one region and few in others one can study the effects across study participants, and with appropriate micro data also the spillovers on non-participants.[21] Since many public employment agencies already work with large populations and online advice is scalable, they seem best placed for such steps. Third, one can fine-tune the type of advice that is provided. Here we focus on coarse occupational advice based on participants' individual

[20] See Section 3.4 for details.

[21] See Crépon et al. (2013) for such a design for traditional unemployment services.

jobgoals. The upside is that occupational advice is easy to learn, and can be used in other situations such as when talking to friends and family. Since many jobs arise from informal job search, such informational spillovers might be crucial. Still, machine-learning algorithms might provide even more individualized advice, even though their suggestions might be less transparent to job seekers and issues of algorithmic bias based on attributes like gender or race are still under investigation in this literature.

References

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

Behaghel, L., Crépon, B., & Gurgand, M. (2014). Private and public provision of counseling to job seekers: Evidence from a large controlled experiment. *American economic journal: applied economics*, *6*(4), 142–74.

Belot, M., Kircher, P., & Muller, P. (2019). Providing advice to jobseekers at low cost: An experimental study on online advice. *The review of economic studies*, *86*(4), 1411–1447.

Bennmarker, H., Grönqvist, E., & Öckert, B. (2009). *Effects of outsourcing employment services: evidence from a randomized experiment* (Tech. Rep.). Working Paper.

Bergemann, D., & Ottaviani, M. (2021). Information markets and nonmarkets. In *Handbook of industrial organization* (Vol. 4, pp. 593–672). Elsevier.

Card, D., Kluve, J., & Weber, A. (2010). Active labour market policy evaluations: A meta-analysis. *The economic journal*, *120*(548), F452–F477.

Card, D., Kluve, J., & Weber, A. (2018, 10). What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations. *Journal of the European Economic Association*, *16*(3), 894-931.

Cottier, L., Fluckiger, Y., Kempeneers, P., & Lalive, R. (2018). Does job search assistance really raise employment? *Available at SSRN 3249880*.

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., & Zamora, P. (2013). Do labor market policies have displacement effects? evidence from a clustered randomized experiment. *The quarterly journal of economics*, *128*(2), 531–580.

Dhia, A. B., Crépon, B., Mbih, E., Paul-Delvaux, L., Picard, B., & Pons, V. (2022). *Can a website bring unemployment down? experimental evidence from france* (Tech. Rep.). National Bureau of Economic Research.

Duell, N., Thurau, L., & Vetter, T. (2016). *Long-term unemployment in the eu: Trends and policies.* Bertelsmann Stiftung Gütersloh.

Gautier, P., Muller, P., van der Klaauw, B., Rosholm, M., & Svarer, M. (2018). Estimating equilibrium effects of job search assistance. *Journal of Labor Economics*, *36*(4), 1073–1125.

Horton, J. J. (2017). The effects of algorithmic labor market recommendations: Evidence from a field experiment. *Journal of Labor Economics*, *35*(2), 345–385.

Krug, G., & Stephan, G. (2013). *Is the contracting-out of intensive placement services more effective than provision by the pes? evidence from a randomized field experiment* (IZA Discussion Papers No. 7403). Institute of Labor Economics (IZA). Retrieved from `https://EconPapers.repec.org/RePEc:iza:izadps:dp7403`

Kuhn, P. J. (2014). The internet as a labor market matchmaker. *IZA World of Labor*.

Layard, R. (2005). Preventing long-term unemployment. In *Working for full employment* (pp. 217–232). Routledge.

Ljungqvist, L., & Sargent, T. J. (2008). Two questions about european unemployment. *Econometrica*, *76*(1), 1–29.

Machin, S., & Manning, A. (1999). The causes and consequences of longterm unemployment in europe. *Handbook of labor economics*, *3*, 3085–3139.

Naya, V., Bied, G., Caillou, P., Crépon, B., Gaillac, C., Pérennes, E., & Sebag, M. (2021). Designing labor market recommender systems: the importance of job seeker preferences and competition. In *4th IDSC of IZA workshop: Matching workers and jobs online-new developments and opportunities for social science and practice.*

Van den Berg, G. J., & Van der Klaauw, B. (2006). Counseling and monitoring of unemployed workers: Theory and evidence from a controlled social experiment. *International economic review*, *47*(3), 895–936.

A Appendix

A.1 Additional empirical results

Figure A.1: Registration and activity over time

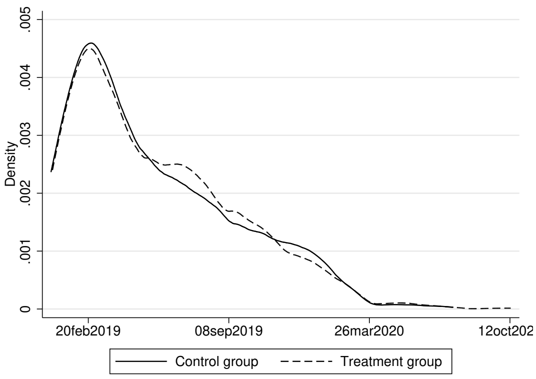

(a) Registration dates

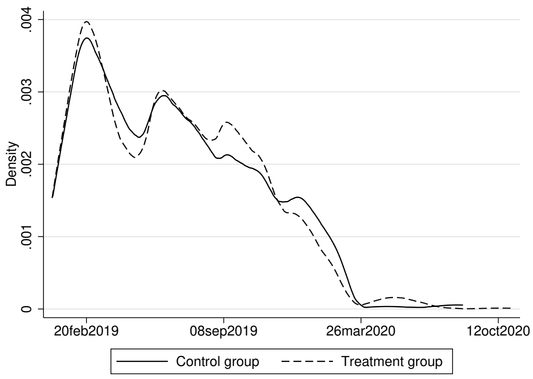

(b) Search activity

Table A.1: Outcomes (job related): only participants that searched at least once on the platform

Variable	N	(1) control Mean/SD	N	(2) treatment Mean/SD	T-test P-value (1)-(2)
Number of interviews	287	0.87 (1.68)	492	0.77 (1.28)	0.34
Duration (days) until first interview	123	139.45 (119.18)	213	129.83 (126.70)	0.49
Job found	298	0.33 (0.47)	502	0.34 (0.47)	0.81
Duration (days) until first stable job	55	167.27 (110.97)	108	129.82 (94.38)	0.03**
Duration (days) until earnings threshold	45	267.69 (111.68)	84	240.70 (117.23)	0.21
Duration (days) until first job	99	156.71 (114.27)	171	131.89 (100.24)	0.06*

***, **, and * indicate significance at the 1, 5, and 10 percent critical level. Jobs that last less than 3 days are omitted. 'N' refers to the number of observations included in the mean/SD. All duration variables exclude individuals that did not secure (within our observation period) an interview/stable job/earnings threshold/first job.

Table A.2: Outcomes (job related)

Variable	N	(1) control Mean/SD	N	(2) treatment Mean/SD	T-test P-value (1)-(2)
Number of interviews	506	0.75	849	0.70	0.48
		(1.44)		(1.18)	
Duration (days) until first interview	206	144.43	346	133.64	0.32
		(120.99)		(122.91)	
Job found	525	0.32	869	0.32	0.93
		(0.47)		(0.47)	
Reach earnings threshold	525	0.15	869	0.15	0.99
		(0.36)		(0.36)	
Duration (days) until first stable job	99	156.17	172	131.26	0.05**
		(106.63)		(95.22)	
Duration (days) until earnings threshold	79	262.33	131	244.08	0.28
		(118.64)		(116.86)	
Duration (days) until first job	168	156.22	280	138.46	0.09*
		(111.59)		(103.32)	

***, **, and * indicate significance at the 1, 5, and 10 percent critical level. Jobs that last less than 3 days are omitted. 'N' refers to the number of observations included in the mean/SD. All duration variables exclude individuals that did not secure (within our observation period) an interview/stable job/earnings threshold/first job.

Table A.3: Fraction of time employed during the first 15 months

	(1)	(2)	(3)	(4)
Treatment group	0.016	0.028	0.043*	0.0070
	(0.016)	(0.021)	(0.024)	(0.039)
Observations	1272	734	412	322
Control variables	Yes	Yes	Yes	Yes
Sample	All	Searched on platf.	Searched on platf. and long unempl.	Searched on platf. and short unempl.

Robust standard errors in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level.

Table A.4: Duration analysis until first stable job/earnings threshold: those that searched on the platform

	Stable job			Earnings threshold		
	(1)	(2)	(3)	(4)	(5)	(6)
	6 months	1 year	End	6 months	1 year	End
Treatment group	1.26	1.23	1.20	1.51	1.12	1.11
	(0.26)	(0.21)	(0.20)	(0.54)	(0.23)	(0.21)
Observations	800	800	800	800	800	800
Control variables	Yes	Yes	Yes	Yes	Yes	Yes

Exponentiated coefficients. Robust standard errors in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level.

Table A.5: Duration analysis until first stable job/earnings threshold: all participants

	Stable job			Earnings threshold		
	(1)	(2)	(3)	(4)	(5)	(6)
	6 months	1 year	End	6 months	1 year	End
Treatment group	1.23	1.12	1.11	1.13	1.03	1.01
	(0.19)	(0.15)	(0.14)	(0.29)	(0.16)	(0.15)
Observations	1394	1394	1394	1394	1394	1394
Control variables	Yes	Yes	Yes	Yes	Yes	Yes

Exponentiated coefficients. Robust standard errors in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level.

Table A.6: Duration analysis until first stable job/earnings threshold: only those searched on the platform and have long unemployment duration (more than 18 months)

	Stable job			Earnings threshold		
	(1)	(2)	(3)	(4)	(5)	(6)
	6 months	1 year	End	6 months	1 year	End
Treatment group	1.88**	1.52*	1.48	.	1.63	1.60
	(0.60)	(0.38)	(0.36)	(.)	(0.51)	(0.46)
Observations	478	478	478	478	478	478
Control variables	Yes	Yes	Yes	Yes	Yes	Yes

Exponentiated coefficients. Robust standard errors in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level. The estimate in column (4) is missing due to an insufficient number of individuals reaching the income threshold within six months.

Table A.7: Duration analysis until first stable job/earnings threshold: only those searched on the platform and have short unemployment duration (less than 18 months)

	Stable job			Earnings threshold		
	(1)	(2)	(3)	(4)	(5)	(6)
	6 months	1 year	End	6 months	1 year	End
Treatment group	1.02	1.15	1.11	0.66	0.85	0.84
	(0.28)	(0.27)	(0.26)	(0.27)	(0.23)	(0.21)
Observations	322	322	322	322	322	322
Control variables	Yes	Yes	Yes	Yes	Yes	Yes

Exponentiated coefficients. Robust standard errors in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level.

Figure A.2: Coefficients linear regressions job finding: first job

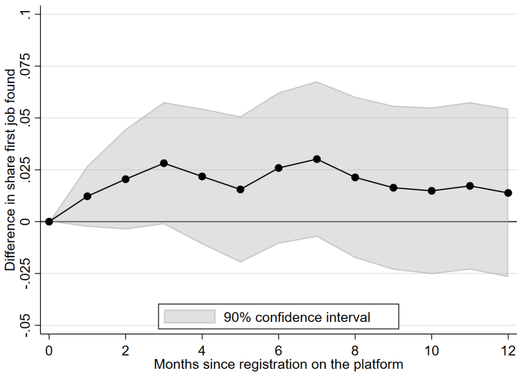

(a) Everyone

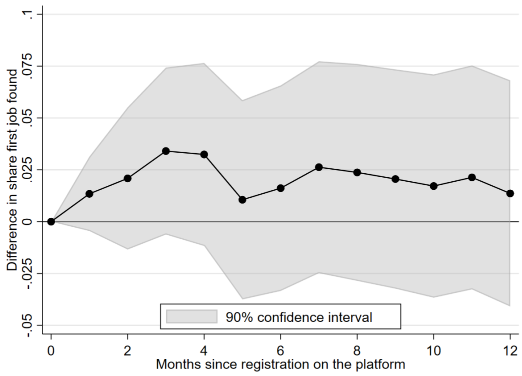

(b) Only those that searched on the platform

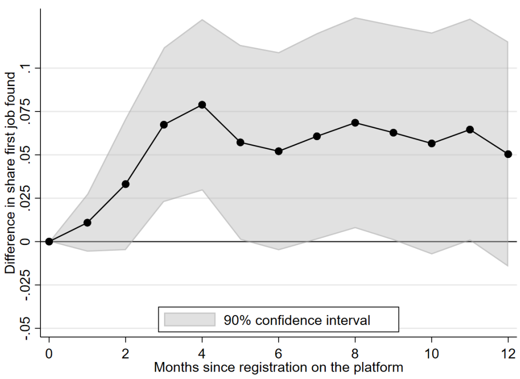

(c) Only those that searched on the platform + long unemployed

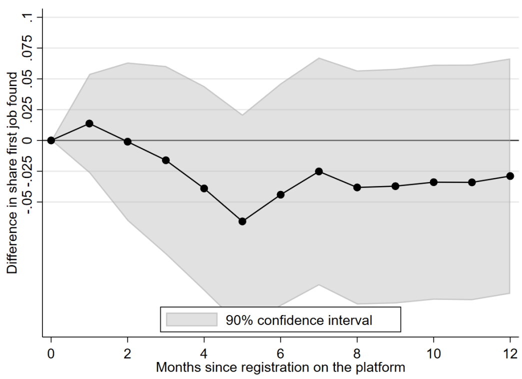

(d) Only those that searched on the platform + short unemployed

Table A.8: Comparison of characteristics of job seekers with long- and short unemployment duration

Variable	(1) Short duration Mean/SD	(2) Long duration Mean/SD	T-test P-value (1)-(2)
Female	0.38	0.41	0.24
	(0.49)	(0.49)	
Non-white ethnicity	0.29	0.23	0.02**
	(0.45)	(0.42)	
Married/Cohabiting	0.11	0.09	0.41
	(0.31)	(0.29)	
Children	0.17	0.16	0.55
	(0.37)	(0.36)	
Jobseeker's allowance	0.27	0.47	0.00***
	(0.45)	(0.50)	
Unemployment duration (months) [a]	12.63	102.50	0.00***
	(6.18)	(50.90)	
Never worked before	0.00	0.15	0.00***
	(0.00)	(0.36)	
Disability	0.54	0.56	0.52
	(0.50)	(0.50)	
N	569	825	

***, **, and * indicate significance at the 1, 5, and 10 percent critical level. Short duration means unemployment duration at the time of entering the program was at most 18 months. Long duration means unemployment duration at the time of entering the program was above 18 months. P-values refer to a test for equality between control and treatment group. [a] Unemployment duration in months at the time of starting on the WHP programme, excluding individuals that never worked before.

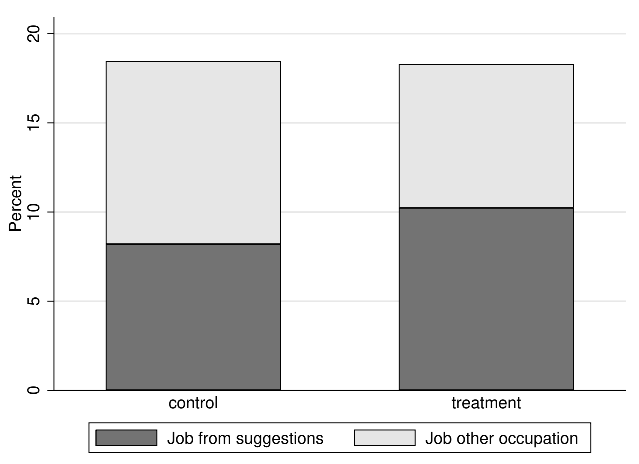

Figure A.3: Occupations of the first stable job found

Figure A.4: Occupational distribution of jobs found

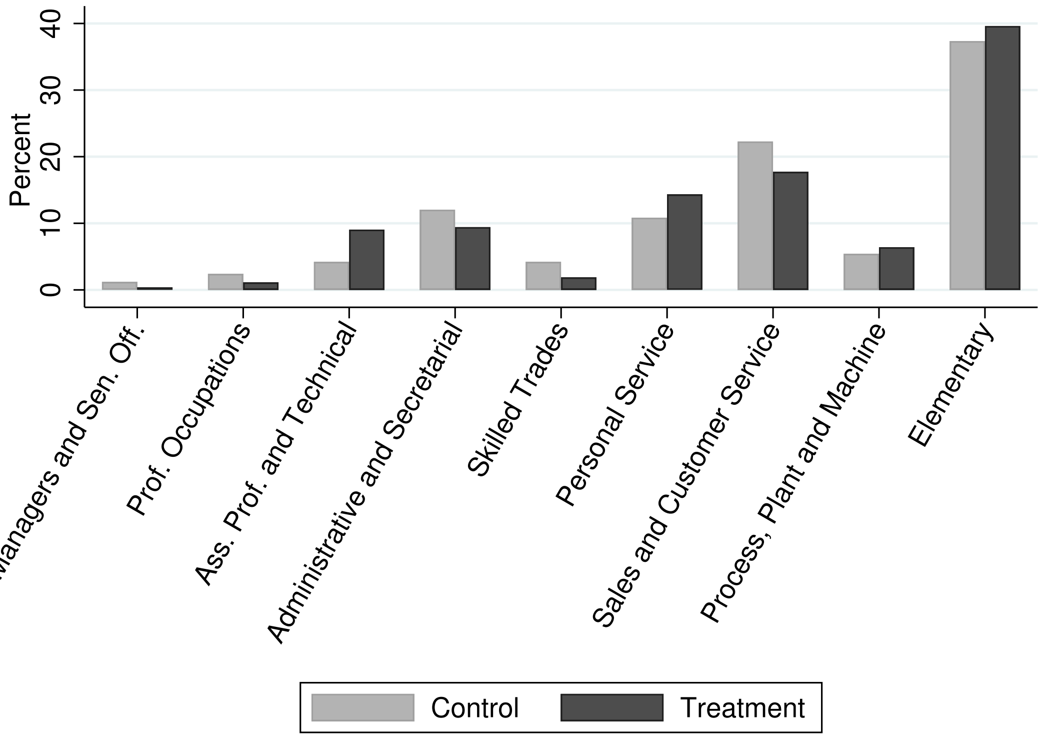

Table A.9: Job search behaviour

Variable	N	(1) control Mean/SD	N	(2) treatment Mean/SD	T-test P-value (1)-(2)
Jobgoals:					
Saved at least one jobgoal	525	0.26	869	0.36	0.00***
		(0.44)		(0.48)	
Mean nr saved jobgoals	525	0.86	869	0.90	0.70
		(2.12)		(1.72)	
Saved at least one manual jobgoal	525	0.26	869	0.25	0.77
		(0.44)		(0.43)	
Mean nr manually saved jobgoals	525	0.86	869	0.63	0.02**
		(2.12)		(1.48)	
Saved at least one no-qual. jobgoal	525	0.00	869	0.16	0.00***
		(0.00)		(0.36)	
Mean nr saved no-qual. jobgoals	525	0.00	869	0.26	0.00***
		(0.00)		(0.77)	
Search queries[a]:					
Total nr of jobgoal searches	525	0.93	869	0.90	0.95
		(5.75)		(5.97)	
Total nr of keyword searches	525	1.75	869	1.72	0.91
		(5.00)		(5.78)	
At least one suggestions-search[b]	525	0.00	869	0.25	0.00***
		(0.00)		(0.44)	
Average nr of included jobgoals (excl sug's)[c]	113	2.97	222	2.17	0.00***
		(2.30)		(1.54)	
Average nr of included suggestions[d]	113	0.00	223	7.79	0.00***
		(0.00)		(5.21)	
Average radius in jobgoal searches	113	8.94	223	13.93	0.43
		(4.11)		(66.90)	
Average radius in keyword searches	234	8.50	382	8.41	0.76
		(3.96)		(3.45)	
Search query results:					
Median nr of results in jobgoal searches[e]	113	1460.65	223	5230.49	0.00***
		(3324.77)		(12997.96)	
Median nr of results in keyword searches[e]	234	9414.19	382	12834.30	0.24
		(28783.16)		(38252.30)	

The sample sizes N refer to the number of individuals that are included in the respective means. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level. [a] All variables are first averaged by individual (across search queries) and then averaged across individuals. [b] Indicator for whether the individual performed at least one jobgoal search with the suggestions included. [c] The average number of jobgoals that were included in the jobgoal search queries (this does not include the suggestions that may have also been included in the search query). [d] The average number of included suggestions within jobgoal searches. [e] The average over the median number of search results per individual.