



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Providing government assistance online: A field experiment on employment assistance

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Providing government assistance online: A field experiment on employment assistance / Giulio Zanella; Guglielmo Briscese. - In: JOURNAL OF POLICY ANALYSIS AND MANAGEMENT. - ISSN 1520-6688. - ELETTRONICO. - 41:2(2022), pp. 579-602. [10.1002/pam.22368]

Availability:

This version is available at: <https://hdl.handle.net/11585/835380> since: 2022-10-25

Published:

DOI: <http://doi.org/10.1002/pam.22368>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Briscese, G., Zanella, G., & Quinn, V. (2022). Providing Government Assistance Online: A Field Experiment with the Unemployed. *Journal of Policy Analysis and Management*, 41(2), 579-602.

The final published version is available online at:

<https://doi.org/10.1002/pam.22368>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Providing government assistance online: Evidence from a field experiment with the unemployed*

Abstract

Welfare programs often consist of mandated in-person assistance services. This feature can introduce an engagement barrier for some beneficiaries. Offering some of these services online can address this problem while also reducing administrative costs. In a field experiment with about 2,700 beneficiaries of unemployment benefits, we evaluate the effectiveness of a self-directed website that supplements assistance traditionally delivered by job center staff. Tracking employment outcomes for nearly two years, we find that the intervention significantly increased job-finding rates for some groups. Towards the end of the first year, the effect is still 7 percentage points (25% higher than in the control group) for prime-age job seekers (35-50 years old) and 9 percentage points (35% higher than the control group) for women, reversing the job-finding gender gap. We discuss opportunities for governments to scale up similar low-cost interventions to assist social insurance and welfare beneficiaries online.

JEL Classification: H53, I38, J64, J68

Keywords: online government services, online behavior, job search assistance

*Ethics approval No. H-2019-232. This study is registered in the AEA RCT Registry.

1 Introduction

Assistance to welfare beneficiaries often consists of labor-intensive services delivered in person by trained staff of a government agency or a delegated provider. These services are costly and involve frequent interactions with caseworkers, which may discourage the participation of beneficiaries for whom attendance is difficult or unattractive. A possible solution is to deliver some assistance online. Remote assistance may increase take up of services by removing the face-to-face constraint and would allow for cheaper customization to better meet a beneficiary's needs. However, the opposite effect is also possible if beneficiaries feel more disconnected than in an offline environment, especially those who are more disadvantaged or not computer savvy. These questions deserve attention: despite the potential benefits of delivering assistance online, there is limited evidence about its effectiveness.

We address this question by evaluating a self-help online tool for the unemployed, a context that is of particular relevance for several reasons: (i) re-employment services are widely adopted by OECD countries; (ii) they are traditionally expensive to implement and difficult to reform, and share key features with many other government services, including a relationship between a case manager and beneficiaries and an outcome framework to measure effectiveness; (iii) re-employment is seen as a poverty alleviation tool and often receives greater attention by policymakers and politicians compared to other welfare programs. The cost-effectiveness of these programs has been questioned (Crépon and Van Den Berg, 2016). Although there is evidence that better targeting can help (Card et al., 2018), little is known about whether changing a service delivery mode can improve targeting. Studies have shown that job seekers can be helped remotely by nudging them to consider vacancies on an online job portal that they might have missed (Belot et al., 2019), but the lack of information is only one of the barriers that job seekers face. Another source of search frictions is the lack of certain job search skills that enable individuals to engage effectively with employers who have vacancies. A job seeker may be informed about suitable job opportunities and yet lack basic skills to pursue them, such as writing an effective resume or an informative cover letter. Submitting low-quality applications may prevent even motivated, well-informed job seekers from obtaining a job (Thoms et al., 1999).

Motivated by these considerations, we evaluate a tool consisting of key, low-cost resources available at a website that supplements the assistance provided by job centers: (i) a resume template and advice on how to customize it; (ii) a cover letter template and advice on how to personalize it; and (iii) tips on how to look and apply for jobs. The evaluation consists of a randomized control trial (RCT) involving about 3,000 job seekers that was implemented in partnership with the then Australian Department of Jobs and Small Business (‘the Department’, henceforth) and a job service provider in Sydney. The outcome of interest in our analysis is whether a job seeker leaves unemployment by finding a job on their own. We estimate the average Intention-To-Treat effect (ITT) employing both simple regression analysis and duration models. Given that our treatment consists of access to different resources, we can only estimate the ITT of the bundle. However, we argue below that bundling these particular treatments is both appropriate and policy relevant.

We find that the intervention significantly sped-up re-employment, with large gains for individuals in the middle of prime age, particularly women: among job seekers in the age range 35–50 the ITT is about 7 percentage points (p.p.) from the third month (off a base of 10.9% in the control group) and persists until the end of the first year (off a base of 27.1%) before slowly fading away as job seekers in the control group also find a job. For women in this age group, the effect is even larger: 9.1 p.p. after three months from assignment to treatment (off a base of 8.1%) and about 10 p.p. from the fourth month to almost the end of the first year (off a base of between 10.8% and 22.1%). Such a gender difference reverts the gender gap in the job-finding rate at the baseline. A simple cost-benefit analysis indicates a gross benefits of between \$1,000 and \$1,350 per job seeker per month of anticipated exit from unemployment, vis-a-vis a very small cost. We also find evidence of a possible displacement of search effort for job seekers older than 50. For them, the point estimates are negative, particularly between the second and fourth month, which may reflect a reduced ability to make a good use of online employment assistance among those who are older and possibly less technologically proficient. Thus, some job seekers could be directed to receiving assistance online by default before accessing intensive counseling at a job center. This would free up resources of frontline staff to provide closer assistance to less digitally savvy job seekers, thereby improving the overall cost-effectiveness of employment services.

The context of employment assistance services provides a compelling case study for the design of welfare assistance programs in general. With more citizens around the world being connected online and digitally savvy, our insights are valuable for the improvement of a broad range of government services. A move towards delivering welfare assistance online presents several advantages. First, both the cost of providing and receiving services would decrease. This would be particularly so for time-constrained citizens, such as those with caring responsibilities, for whom turning up to regular mandated appointments can be costly and who would also prefer to access some of these services at their own pace. The online environment also offers a wide range of already freely available training platforms (e.g. MOOC or Coursera) that government services can help match to citizens based on needs. Second, such environment allows for program targeting at a lower marginal cost. Policymakers can monitor take up and usage of online resources, whether it being enrolling in a program or attending an online training session, and improve their services at a lower cost. Third, the online environment allows for faster and more accurate evaluation of what works and for whom. Assisting beneficiaries online would enable policymakers to (i) track intermediate outcomes, such as the level of engagement with training materials, better than in an offline world; (ii) measure job seekers' preferences and satisfaction with the assistance they receive, in real-time and at virtually no cost. Our results also show that not everyone may benefit equally from an online service delivery mode. For these beneficiaries, support could continue to be provided face-to-face, or even at a more intensive level thanks to the resources that are freed up. Experiments like ours are pivotal in determining not just what works, but also for whom and how, so to ensure that services are not rolled out universally if they are not effective for some groups.

Our findings also contribute to the literature on improving active labor market policies. Reviewed in more detail below, a limited number of field studies have explored the provision of light-touch employment assistance on labor market outcomes. These include two studies demonstrating that employment outcomes can be enhanced by providing brief information to older workers (Liebman and Luttmer, 2015) or informative brochures to job seekers who recently lost their job and are at risk of long-term unemployment (Altmann et al., 2018). In the digital context, a low-cost intervention adding job search advice in a job search platform

led to an increased number of job interviews (Belot et al., 2019). However, to date no field studies have assessed the impact of providing a comprehensive package of job assistance strategies and resources online. The potential of digital technologies was also studied by Stevenson (2009) and Kuhn and Mansour (2014), who document that job seekers using the Internet to find employment are more likely employed. More recently, Gürtzgen et al. (2021) find that access to high-speed Internet in Germany improves reemployment rates early in the unemployment spell. Our results suggest that these tools would be even more effective when complemented by resources that supplement job search skills.

The rest of the paper is organized as follows. Section 2 describes the institutional context in which the experiment was implemented, the pre-intervention study that we conducted to inform the design of the online resources, and our randomization. The results are presented in Section 3, and Section 4 concludes.

2 The experimental design and its implementation

2.1 Institutional background

The active labor market policies (ALMP) system in Australia is called ‘Jobactive’ and consists of contracts awarded by the Commonwealth Department of Employment to providers that deliver re-employment services to job seekers and that receive a performance-based compensation. The level of job readiness is first assessed by a federal government agency, which is then responsible for matching unemployed workers with job service providers. The Department of Employment can therefore monitor the entry and exit of job seekers in the system and measure providers’ performance. In many ways, the Australian system is closer to the UK and other European systems than to the US, where ALMPs are decentralized at the state level. Jobactive also includes a compliance framework designed to encourage job seekers to engage with their provider, undertake activities to meet their mutual obligation requirements, and demonstrate that they are actively looking for work. One of these activities is regular appointments (usually every other week) with a job advisor at their registered provider’s center. During their first meeting, job seekers complete a job plan, and the job

advisor can determine which programs or services they might be eligible for and benefit from. In subsequent meetings, job advisors can help a job seeker improve their resume and cover letter or enroll them in a training course to learn how to do so. While job seekers must be “work-ready” to enrol with a job service provider, the level of readiness varies. Thus, some unemployed workers are able to find a job with minimal assistance while others require more guidance. This insight inspired the design of light-touch and easy-to-use resources to help job seekers help themselves improve the quality of job applications.¹

2.2 Pre-intervention study

These easy-to-use resources were designed based on a combination of best-practice job search techniques – identified through a literature review, as detailed below – and qualitative formative research that we undertook before the trial. Our formative research consisted of: (i) a detailed review of Department desktop resources; (ii) a focus group with ten experienced job coaches from our partner provider; (iii) interviews with ten job seekers enrolled with our partner.² Such casual observation suggested that many job seekers were submitting low quality job applications, for example resumes that were not updated or tailored to the vacancy they were applying for, or not attaching cover letters even when required. A variety of reasons explain this behavior, including low awareness of the requirements of job applications, poor feedback from previous applications, perceptions that writing more would make a negative impression due to low literacy levels, and low motivation levels after repeated failures to secure an interview. These suggestions were backed up by the Department’s employer survey data, which indicate that one in five employers reported not interviewing applicants because their written application was not tailored to the position, failed to address selection criteria or had spelling mistakes (DJSB, 2015). In the same survey, employers reported

¹More information about Jobactive is available at <https://www.employment.gov.au/jobactive>. As part of the policies detailed in the compliance framework, job seekers must be actively looking for work to keep receiving their unemployment benefits. The benefit amount depends on their socio-economic situation, health conditions, and family composition, among other factors. For a full list of offered benefits, see: <https://www.servicesaustralia.gov.au/sites/default/files/2017/03/co029-1703.pdf>. For more information about similarities between the Australian and other systems, see instead OECD (2012).

²In order to avoid disruption, we acted as mere observers of job seekers–staff interactions during their regular appointments, and we talked informally to some of these job seekers after these meetings. Due to privacy reasons, we were not able to assess the materials that job seekers used to apply for jobs, which is why we complemented it with a desktop research.

placing a large emphasis on the applicants' communication skills, which one's application package typically reveal. Details such as attaching a cover letter, reporting in one's resume up-to-date and clear professional contact details, work experience, job responsibilities, qualifications and certifications, achievements, and professional objectives have been shown to be important across a number of industries (Hornsby and Smith, 1995). In addition to grammatical errors, certain pieces of information are instead better avoided, such as high school education – unless it is the most recent – or demographic details such as age, marital status, and parental status (Schramm and Dortch, 1991; Toth, 1993; Thoms et al., 1999; Ross and Young, 2005). Since one's resume and cover letter are the first opportunity to demonstrate communication skills, job seekers who are unaware of these guidelines miss out on leaving a good first impression and send a negative signal about their qualities to a potential employer.

A similar pattern emerged for job search strategies. Some job seekers did not use free online job search platforms, some appeared to inefficiently trade-off application quality for quantity, and others applied repeatedly to the same employer despite previous rejections, not expanding their job search geographically or by industry. While it is part of the responsibilities of job advisors to support job seekers with job search resources and strategies, it is often difficult to provide tailored assistance due to multiple job seekers requiring attention at the same time and other administrative responsibilities. Moreover, job advisors can persuade job seekers to improve their job search behavior, but close monitoring of every application for every assisted person is not feasible.


2.3 Intervention resources

The picture emerging from this preliminary study was thus one of poor job search skills, possibly for a substantial fraction of job seekers. To supplement these skills, we designed a resume and cover letter templates that contain all the right features described in Section 2.2 and we compiled evidence-based job search strategies into a short checklist containing ten tips to follow when pursuing vacancies. We then set up to make these three key job search resources available online, on a purposely designed website that was called 'My Job Goals' (MJG) and that would be accessible both at job centers and at job seekers' homes. These resources are reproduced in Sec.1 of the online appendix. The website was developed in-house

at the cost of about 40 hours of researchers and staff time (alternatively, its development could have been outsourced at a modest cost) and it required a mere \$25 a month for maintenance, so our intervention is both scalable and easy to replicate in other contexts.

Since we jointly analyze the effect of a “bundle” of treatments – resume and cover letter guidance, and tips for improving the search process – we cannot disentangle which specific treatment worked and which, if any, did not. However, bundling these three treatments is both appropriate and policy relevant. It is appropriate because they are complementary, which implies that evaluating each of them in isolation rather than in a bundle may be misleading – for example, a high-quality resume may be ineffective if attached to a low-quality cover letter, and vice versa; similarly, the effectiveness of high-quality resume and cover letter would be dampened by ineffective search strategies. It is policy relevant because given their cost-effectiveness and complementarity, it would not be wise for the employment agency to offer separately the three assistance tools that we study. As illustrated in Figure 1, MJG was designed to be straightforward and self-explanatory:

Figure 1: The ‘My Job Goals’ website



Notes: Left: the homepage of the ‘My Job Goals’ (MJG) website, an online resource designed for the experiment and containing editable resume and cover letter templates and a checklist with tips on how to look and apply for jobs. The website was accessible to job seekers both at job centers and at home between May and August 2017. The logo of the provider was clearly visible at the top-left of the screen on each page, but was removed from this picture for confidentiality reasons. Right: how ‘My Job Goals’ (MJG) was accessible at a job center. A posters advertising the website is also visible in this picture.

the homepage had clear and visible links to three webpages, each focused on one of the resources described above. Each webpage contained a brief explanation on the importance of the resource for job seeking, and a downloadable template in MS Word for the resume and cover letter. To reinforce further the value of the webpage content, each page had a short video explanation from an experienced staff member of the job service provider. The

style of the website, including colors, photos, and fonts, purposely resembled the marketing and communication material of the job service provider. The logo of the provider was also clearly visible at the top of the homepage and each webpage to ensure familiarity and trust with the source of the communication. [Figure 1](#) also shows how MJG was accessible at one of the treated job centers' search facilities. Posters advertising the website are also visible in this picture. The rationale for providing these resources online was threefold: first, we wanted to test solutions that were replicable and scalable at low-cost; second, relax the time constraint of frontline staff at job centers; third, help job seekers to focus on their online job search and application process.

2.4 Randomized trial design

In partnership with the Department, we sent a call for expression of interest to all Jobactive providers.³ Among those that responded, we selected a provider that had the capacity to work with the research team for a randomized trial in Sydney (NSW) and had enough job centers to allow for site-level randomization. This provider was not randomly selected, but this has little consequences for our analysis. Internal validity is certainly unaffected. As for external validity, we purposely designed an intervention that required minimal input by provider's staff, so that the provider's identity has little relevance. Although the RCT was carried out in a metropolitan area, the treatment was specifically designed so that it could be scalable and replicable also in harder-to-reach locations across the country.

A field experiment that randomized exposure to MJG was then conducted between May 15 and August 15, 2017, a time when the Australian labor market was characterized by a relatively low unemployment rate (5.6% in June 2017) and a labor force participation rate that had remained fairly stable at about 65% for 10 years. In order to minimize spillovers onto job seekers in the control group via shared treatment material and resources, the randomization was conducted at the job center level. The selected providers had 22 job centers in Sydney, and sites with shared job advisors were pooled. Given the small number of sites,

³There are many providers in the system, which vary in terms of performance. The Jobactive star rating system is a tool to evaluate providers' performance in terms of placement and job retention. The Department issues a report every quarter. All reports are available at www.employment.gov.au/jobactive-star-ratings-and-performance.

in order to keep chance imbalances between treatment groups to a minimum and to achieve better estimates of the ITT, we employed a rerandomization procedure (Morgan and Rubin, 2012).⁴ This procedure involved randomly allocating sites to control and treatment 1,000 times, and then calculating the relative imbalance on each of the site-level balancing covariates.⁵ Relative imbalance was calculated by taking the absolute value of the t -stat from a simple regression with treatment allocation as the outcome variable and the balancing variable as the sole predictor. We then took the largest t -stat of the balance checks for each allocation, and we chose the allocation that had the smallest maximum imbalance. The 22 job centers were thus split into 11 treatment sites and 11 control sites. This design is such that both the sampling process and the assignment mechanism are clustered at the job center level, which requires adjusting standard errors appropriately (Abadie et al., 2017). We return on this issue below. Figure 2 shows the spatial location of treated and control centers in the Sydney metropolitan area that was produced by this randomization procedure. The counterfactual in our experiment is “business as usual”, i.e., control job centers continued their normal activity and offered the same programs to job seekers. Moreover, during the trial period all the control or treatment centers did not implement any new training courses.

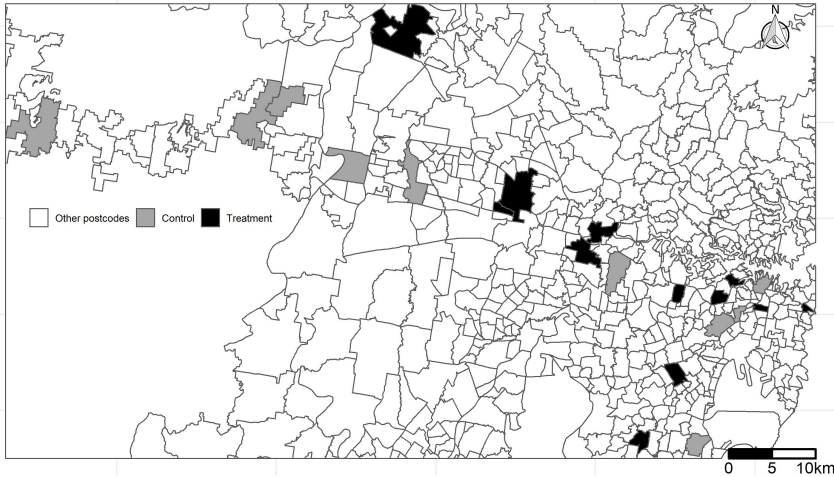
Particularly important in a cluster-randomized field intervention (Rose and Bowen, 2009), our experiment was well-powered. Taking the outcome of the randomization as given, we calculated the power to detect a range of effect sizes, thus also determining the duration of the trial that we needed. Specifically, we calculated power by taking the number of observations in the control and treatment groups and deflating them by the design effect for a cluster-randomized trial with unequal cluster size. This design effect was given by $\frac{\bar{n}c}{\sum_{i=1}^c \frac{n_i}{1+(n_i-1)\rho}}$, where c is the number of clusters, n_i is the number of individuals in each cluster, \bar{n} is the average number of individuals across clusters, and ρ is the intracluster correlation coefficient (ICC; Rutterford et al., 2015; Kerry and Martin Bland, 2001). We first estimated the ICC

⁴Randomization using stratification was not possible due to the high number of balance variables and the low number of sites.

⁵There covariates are: part-time or full-time status of the site; East/ West status of the site (to account for socioeconomic disparities in Greater Sydney along East and West divisions); caseload of wage subsidy eligible job seekers; caseload of job seekers at the site; single or composite site (as mentioned, due to some crossover in job advisor support to different sites, we created composite sites so that any advisors would work in either a control or a treatment center); average monthly job placements at the site. These variables are summarized in Panel B of Table 1.

and design effect in our baseline data, and then we used these parameters to calculate the power to detect a given effect size. For total placements in one month, we used as our outcome variable the likelihood of being placed into employment across all job seekers, based on historical data that indicated that approximately 8% of the total caseload at each site was placed into employment each month. We calculated the minimum detectable effect sizes given a power of 80%, for each possible duration of the trial from one to twelve months. This calculation revealed that in three months we were powered to detect a 3.5 percentage-point (or higher) increase in placements. Such an effect size is in line with [Altmann et al. \(2018\)](#) – who find that after 10 months since treatment the job-finding rate among treated job seekers at risk of long-term unemployment is some 2.5 percentage points higher than among control ones – and [Belot et al. \(2019\)](#) – who find that after 3 months into treatment the analogous effect is about 6 percentage points.⁶

Figure 2: Spatial location of job centers assigned to treatment or to the control group



Notes: The figure shows a map of postcode areas in the metropolitan area of Sydney, NSW. The postcodes containing the 11 job centers assigned to treatment and the 11 job centers acting as control centers are filled in black and in gray, respectively.

In our baseline analysis, the treatment group is given by job seekers who were registered at the provider’s treated job centers at the beginning of the trial, a total of 1,442 persons. The control group is given by 1,248 job seekers who, at that same date, were registered at

⁶The power calculation was conducted in the full sample and not also in the smaller subsamples used for the analysis of treatment effect heterogeneity. Given the smaller sample sizes, the analysis in these subsamples is likely underpowered, which leads to more conservative estimates of statistical significance.

the provider’s job centers that were assigned to control. In this baseline analysis we do not consider 1,039 job seekers who registered at the provider’s job centers during the trial (541 at treated job centers and 498 at control ones) because these “latecomers” have been exposed to MJG for a shorter period of time, possibly just a few days.⁷

Table 1 summarizes observable pre-treatment characteristics by treatment status. Panel A contains job seeker characteristics and Panel B contains job center characteristics. The randomization balanced virtually all of these, thus boosting the credibility of the control group as a valid counterfactual for treated job seekers.⁸ The sampling design implies that our sample was representative of the population of job seekers in Sydney at the time the trial was implemented. However, since that was a time when national and state unemployment rates were low, our sample contains job seekers who may have been particularly hard to re-employ. Note that the duration of the current unemployment spell is not observed in our data. What is observed and is reported in Table 1, is the date a job seeker first entered the Jobactive system, i.e., when the *first* recorded unemployment spell began. On average, this is nearly five years before the trial started, for both treated and control job seekers. Since then, one can leave and re-enter the unemployment state multiple times. Given that all job seekers in our data set were unemployed at the time the trial started and that the date one first entered the Jobactive system is balanced, the unobserved duration of the current unemployment spell is presumably also balanced.

Treated job seekers were directed to the website via digital and in-person channels. The digital channel consisted of text and email messages, sent on the trial launch date to the job seeker’s mobile number and email address.⁹

⁷Figure A-4 in the online appendix shows that our results are robust to including these “latecomers” in our estimation samples.

⁸The p -value from a joint significance test – obtained from an F test of the null hypothesis that the coefficients from a regression of the treatment indicator on the covariates in Table 1 are jointly zero – is 0.67 at the job seeker level and 0.92 at the job center level. As usual, the possible effects of the few imbalances in a finite sample can be eliminated by using the individual-level, pre-treatment covariates as conditioning variables when regressing the outcome on the treatment indicator.

⁹Valid phone numbers were available in the Department’s database for 85% of treated job seekers but, as explained in more detail in Sec.2 of the online appendix, only 26% of them had a valid email addresses.

Table 1: Characteristics of job seekers and job centers in treatment and control groups

	Controls		Treated		<i>p</i> -value
<u>Panel A: Job seekers</u>					
Female	0.497	(0.500)	0.469	(0.499)	0.147
Age	39.98	(13.00)	39.89	(12.77)	0.862
College degree	0.142	(0.349)	0.144	(0.351)	0.899
Vocational degree	0.254	(0.435)	0.252	(0.435)	0.925
High School degree	0.238	(0.426)	0.241	(0.428)	0.872
Less than high school	0.240	(0.427)	0.243	(0.429)	0.850
Missing education information	0.010	(0.098)	0.004	(0.064)	0.084
Aboriginal/Torres Strait Islander	0.076	(0.265)	0.077	(0.267)	0.934
People with disability	0.325	(0.468)	0.326	(0.469)	0.938
Primary caregiver	0.127	(0.334)	0.135	(0.342)	0.550
Homeless	0.089	(0.286)	0.109	(0.312)	0.099
Ex-offender	0.097	(0.296)	0.113	(0.317)	0.176
Worked in pre-intervention year	0.260	(0.439)	0.271	(0.445)	0.500
Months since enrolled in Jobactive	56.46	(34.47)	57.51	(35.23)	0.441
Wage subsidy-eligible	0.558	(0.497)	0.560	(0.497)	0.924
Used wage subsidy (if eligible)	0.019	(0.005)	0.019	(0.005)	0.990
Mobile phone number available	0.947	(0.222)	0.952	(0.212)	0.557
E-mail address available	0.385	(0.487)	0.357	(0.479)	0.130
<i>N</i>	1,248		1,442		
<u>Panel B: Job centers</u>					
Has full-time job advisors	0.778	(0.147)	0.875	(0.125)	0.626
Located in East Sydney	0.556	(0.176)	0.500	(0.189)	0.832
Composite center	0.222	(0.147)	0.250	(0.164)	0.901
Caseload, total	135.3	(33.14)	135.5	(18.98)	0.997
Caseload, wage subsidy-eligible	75.0	(19.22)	74.1	(9.71)	0.969
Monthly placements, total	11.156	(2.073)	10.060	(1.512)	0.682
Monthly placements, wage subsidy-eligible	7.156	(1.090)	6.600	(1.010)	0.716
<i>N</i>	8		9		

Notes: Mean and, in parentheses, standard deviation of observable pre-treatment characteristics of job seekers and job centers in treatment and control groups. The *p*-value refers to a t-test of the null hypothesis that the means are equal across the two groups. Sample: 2,690 workers who were unemployed and registered at the provider's job 22 centers at the start of the trial. There are only 17 job centers in the table because of four composite sites: three aggregating two job centers and one aggregating three.

We don't know how many job seekers opened the website from the link reported in the SMS (although the website metrics reported in Figure A-3 of the online appendix provide an indirect measure) but we know that 32% of those who received the email opened it and, of these, about 25% clicked on the link to the website, on top of any other untraceable visit, such as those made by one of the PCs of the job centers.

The in-person channel consisted of cues in job centers, including posters that we designed and hung in visible spots on the walls of treated job centers (see Figure 1), and communication by job advisors. We visited the treated sites a few days before the trial launch date to instruct job advisors on how to explain the benefits of the websites to job seekers during their regular appointments. To facilitate this work, we provided the job center staff with a one-page document that they could consult if needed. During these visits we also hung the posters and a strip of paper underneath the job center PC screen with the website URL. The posters and job advisors' cheat sheet are reproduced in Figures A-1 and A-2 of the online appendix. Potentially, the in-person channel introduces two additional treatments to the bundle that comes with the website: (i) caseworkers at treated job centers are aware of the trial and so may alter their effort; and (ii) some job seekers are eligible for a wage subsidy and caseworkers at treated job centers are instructed to remind them of this opportunity. To rule out that any of our results are driven by employers making use of wage subsidies, note in Table 1 that among eligible job seekers, the subsidy take-up rate is very small and perfectly balanced across treated and controls. As for caseworker's effort, we show below that our results are similar when two alternative job-finding measures are employed: a "self-found" job or a "provider-sourced" job. The latter reflects caseworker's effort to a greater extent than the former.

On the trial end date we shut down the MJG website and physically removed the posters and any other material from the treatment sites. Website metrics, reported in Section 4 of the online appendix, indicate that MJG was visited 1,659 times during the trial. Traffic peaked when SMS and email messages were sent out. The resume and cover letter template pages were visited by 424 and 260 unique individuals, respectively, for about 2 minutes each on average. This is enough time to inspect the document and download it. The job search tips page was visited by 359 unique visitors, for 2.6 minutes on average. Due to

privacy regulations, we could not track the identity of job seekers who visited the website or downloaded material. As a consequence, we can only identify the intention-to-treat (ITT) effect, i.e., the effect on job-finding of being invited to use MJG via the digital and in-person channels. This is the estimand that we target. Note that it is also the relevant parameter from a policy perspective given that the employment agency cannot force job seekers to use online resources that are made available to them.

Yet, even if we randomized assignment to treatment, identifying the ITT is tricky given that we cannot exclude contamination between treatment and control groups during the trial, i.e., that the Stable Unit Treatment Value Assumption (SUTVA) is violated. There are two main sources of possible violation in our experiment, each with different consequences. First, information spillovers that are specific to our design. The website was freely accessible on the Internet and so job seekers in the control group may have accessed it. For example, they may have heard about it from other job seekers or they may have seen the posters by walking into job centers assigned to the treatment group; or job advisors at treated centers may have appreciated the website resources and they may have shared them with their colleagues at control centers. Thus, some subjects not assigned to treatment may end up being treated and yet they would be part of the control group when estimating the ITT. The spatial distribution of treatment and control centers shown in [Figure 2](#) makes this type of contamination a second-order concern, but even if it were first-order, in our context this is the only possible form of noncompliance under the assumption that treated job seekers do not ignore the information that they receive. Therefore, although SUTVA must hold in order for the ITT to be identified ([Angrist, Imbens, and Rubin, 1996](#)), in our experiment this type of violation results, at worst, into attenuation bias in the estimated ITT. Second, classical search externalities, whereby the job-finding prospects of the control group worsen when treated job seekers become more likely to find employment, because the two groups compete for a pool of vacancies that is given in the short run. This is a possibility in *any* experiment that alters search effort or the search technology of some job seekers, and one that leads to overestimating the treatment effect. However, there are reasons to believe that such bias is negligible in our context. Ours is a small trial relative to the number of vacancies in the Sydney metropolitan area, which makes search externalities less likely.

Moreover, as explained in what follows, our re-employment measure is whether a job seeker found her/his own employment directly, which makes it unlikely that a treated job seeker “poaches” vacancies of employers registered with the same provider.

3 Results

The outcome of interest is whether a job seeker left the unemployment state and became employed at least once by a certain date since the beginning of the trial, which is observed from administrative data maintained by the Department for nearly two years, at a monthly frequency.¹⁰ We consider here job placements secured by the job seeker directly (which we label “self-found”), i.e., not from a vacancy originally secured by the job service provider. This notion of job finding is of primary interest in our evaluation because the aim of MJG is precisely to help job seekers help themselves without a direct, specific intervention of the job center in the finalization of the placement. As remarked above, by defining the outcome this way we also minimize the incidence of search externalities in our experiment. Following a placement, it is a job advisor’s responsibility to register it as self-found or provider-sourced by simply ticking a box. This does not affect the compensation or evaluation of either the provider or the job advisor because a job seeker may have of course benefited from the broad support provided by the job center in finding her or his own job. About 90% of job matches in our data are of the self-found type and Figure A-5 of the online appendix shows that our results are robust to considering a broader notion of job finding, i.e., to defining the

¹⁰At the time of the trial implementation, all providers were required to register job seekers and track their job finding outcomes in a shared and centralized database overseen and managed by the Department. Further, the centralized database constitutes also the basis for the provider performance rating and their related outcome payments. As such, data storage is standardized, and the data is updated in real-time. It is possible that a job match is realized and subsequently destroyed. The employment state of a job seeker who left Jobactive is not directly observable in the Department data. Ideally, in order to understand how treatment shifted the entire hazard function, one should use time since the current unemployment spell – rather than the trial – began. As remarked above, such information is not observed in our data. However, given that the time elapsed since a job seeker first entered the Jobactive system, the duration of the current unemployment spell is presumably also balanced. It follows that given that our sample is representative of job seekers in Sydney, our results have external validity even if we cannot incorporate elapsed unemployment duration in the analysis. Note that the share of job seekers looking for work in regional versus urban metropolitan areas in Sydney is likely very similar to other cities in Australia (e.g., Melbourne, Brisbane), and the characteristics of Australian cities, including job seekers’ proximity to job opportunities, resemble that of many locations in other countries.

employment outcome by also considering matches from a vacancy secured by the provider.

It is well known that the effects of job search assistance are heterogeneous along age and gender dimensions (Card et al., 2010, 2018). In order to explore such heterogeneity in the context of our intervention, we conduct the analysis in the full sample and in seven subsamples resulting from different combinations of job seeker’s age and gender: age 35–50, age 34 and younger, age 51 and older, men, women, men 35–50 years of age, and women 35–50 years of age. We report in Section 5 of the online appendix the analog of Panel A of Table 1 for these seven subsamples – Panel B would be identical across them. These tables show that, with few exceptions, the randomization also balanced the characteristics within the smaller groups, although the null hypothesis that the means of covariates by treatment status are equal is harder to reject in these smaller samples.

3.1 Regression analysis: baseline results

The first step of our causal analysis is OLS estimation of linear probability models that are easy to understand and that impose minimal assumptions. The estimating equation is

$$Y_{im} = \alpha_m + \beta_m D_i + \gamma_m X_i + \epsilon_{im}, \quad (1)$$


where Y_{im} is a binary variable indicating whether, by month m into the trial, individual i obtained a job match of the self-found type, D_i is a dummy indicating whether at the start of the trial job seeker i is registered with a job center assigned to treatment, and vector X_i contains the covariates summarized in Table 1. Therefore, parameter β_m identifies the average ITT after m months, which we estimate for each month $m = 1, 2, \dots, 22$ since the onset of the trial. To guard against the perils of multiple hypothesis testing, we report in Table A-9 of the online appendix the adjusted p -values recommended by Romano and Wolf (2005). We also have data on any job matches in the year preceding the intervention, which we use to estimate equation (1) also for lagged $m = -1, -2, \dots, -12$. These estimates provide a placebo test because the ITT should be zero at all these pre-trial lags. This empirical methodology produces 34 point estimates of β_m and their confidence intervals. Note that our experimental design features clustering at the level of job centers both in the

sampling process and in the assignment mechanism, which requires to adjust standard errors for clustering of unobserved shocks at this same level (Abadie et al., 2017). However, with only 22 job centers in the sample we are subject to the perils of clustering in the presence of few clusters (Cameron and Miller, 2015). We resolve this tension by, first, applying the wild cluster bootstrap (Cameron et al., 2008) and then conservatively selecting the confidence interval associated with the largest between the p -value from a test of $H_0 : \beta_m = 0$ produced by the bootstrapping procedure (which accounts for clustering at the job center level) and the conventional one (which is robust to general heteroskedasticity only).

Figure 3 illustrates such points estimates and their 95% confidence intervals (Table A-10 in the online appendix contains the associated tables and p -values). In the pooled sample of all job seekers (top-left panel), exposure to MJG causes an increase in the job-finding rate of 3.5 percentage points in the fourth month that is marginally significant (p -value = 0.053). Subsequently, job seekers in the control group slowly “catch up” with their treated counterparts and so this effect becomes smaller and then vanishes after about eighteen months.¹¹ As expected, all point estimates are zero in the year preceding treatment. This small average effect masks important heterogeneity, as revealed by the remaining panels of Figure 3. First, the effect is much larger for job seekers between 35 and 50 years of age. For them, exposure to the website induces a job-finding rate that by the third month significantly exceeds the one for the control group by 6.9 percentage points. This effect increases to between 7 and 8 percentage points between the fourth and the eleventh month (after eleven months, the average ITT is still 7 percentage points, which is 25% of the job-finding rate in the control group) and then converges to a statistically insignificant 4 percentage points. The difference between the average ITT for the age group 35–50 and the one estimated in the full sample is statistically significant. The finding that the effect is driven by job seekers in the age range 35–50 is consistent with Card et al. (2018), who report smaller average effects of ALMP for older workers and youth.

¹¹By “catch up” we mean the following. Eventually, that is in hypothetically infinite time, every job seeker obtains a job. What exposure to MJG is supposed to do is to improve one’s search technology, which allows job seekers in the treatment group to find a job *faster*. Such an advantage necessarily fades away as job seekers in the control group are also matched.

Figure 3: ITT effects on job finding, linear probability models



Notes: The figure shows, for each month since entering the trial (positive values on the x-axis) and for the 12 months preceding the trial (negative values on the x-axis), the estimated average intention-to-treat effect (parameter β_m in equation 1), represented by a circle. The outcome Y_{im} is a dummy taking value 1 if by month m into the trial (values on the x-axis) jobseeker i obtained *any* job in a way classified by the job center as “Found Own Employment” (self-found) and 0 otherwise. The figure also shows the 95% confidence interval associated with the largest between the p -value from a test of $H_0 : \beta_m = 0$ produced by the wild cluster bootstrap (clustering at the job center level) and the p -value implied by heteroskedasticity-robust standard errors. Sample: 2,690 workers (964 in age range 35–50; 1,045 of age 34 and younger; 681 of age 51 and older; 1,394 men; 1,296 women; 480 men in age range 35–50; 484 women in age range 35–50) who were unemployed and registered at the provider’s job centers at the start of the trial.

Second, in this age group 35–50, the average ITT is substantially larger for female job seekers than for males. For men the average effect is about 4–5 percentage points in the first year – although never statistically significant – converging to about 3 percentage points thereafter. For women, instead, this effect is twice as large and significant, 9–10 percentage points from the third month and throughout the first year from the intervention, slowly converging to an insignificant 6 percentage points afterwards. This is a large effect considering that, between three months and one year from entering the trial, the job-finding rate for women 35 to 50 years of age in the control group ranges between about 8% and 24% and that it is just above 32% after 22 months. After a year, the average ITT estimated for these women is still 8.8 percentage points, which is 36% of the job-finding rate in the control group. Such gender heterogeneity *reverses* the gender gap in the job-finding rate absent the intervention: for men aged 35–50, the job-finding rate in the control group between three months and one year from entering the trial ranges between 14% and 30%, and it is nearly 37% at the end of the observation period. The finding that exposure to MJG is particularly beneficial to female job-seekers is also consistent with [Card et al. \(2018\)](#), who report larger impacts of ALMP for women. A possible explanation in our context is that the resources available on the website helped women improve the quality of their job applications by making their package more informative to employers. For example, in a set of experiments, [Exley and Kessler \(2019\)](#) find significant differences in self-promotion between men and women. These authors suggest that women have internalized that self-promotion is inappropriate or risky. The standardized job search material provided via MJG may have attenuated the gender gap in signaling skills to employers. Another possible explanation is that women are more receptive than men to the resources made available by the experimenters (i.e., communication material and job coaches’ advice about MJG, in our case), as suggested in different contexts by [Eckel and Grossman \(2008\)](#) and [Angrist et al. \(2009\)](#).

Third, the ITT may be temporarily negative for unemployed workers who are older than 50. For them, exposure to the website reduced the probability of finding a job by 4.2 percentage points by the third month from assignment to treatment (p -value = 0.079). Afterwards, this point estimates slowly reverts towards zero. A possible explanation is that the intervention induces older workers to use new job search tools and methods through


a digital platform they may have a harder time becoming familiar with.¹² However, this interpretation should be taken with a grain of salt given that, as shown in Table A-4 of the online appendix, there is a higher incidence of two notable disadvantaged groups (the homeless and the ex-offenders) in the treatment group than in the control group. This pattern is robust to using the broadest notion of job finding available in our data, namely whether there is a placement regardless of whether it is self-found or not. As shown in Figure A-5 of the online appendix, in this case the confidence intervals are larger but the pattern produced by the point estimates is unchanged. Finally, if we conservatively consider the Romano and Wolf (2005) p -values adjusted for multiple hypothesis testing (Table A-9 in the online appendix), we can still detect an ITT that is statistically significant at the 10% level for job seekers in the age range 35–50 and women in this group throughout the first year.

3.2 Duration analysis

We complement the regression analysis by estimating hazard models, a more common specification in the unemployment literature. Assume the following piecewise-constant proportional hazard of exit from unemployment for job seeker i during day d : $\alpha(d; D_i) = \alpha_d \exp(\beta_d D_i)$, where α_d is the baseline hazard during day d . First, we construct unconditional Kaplan-Meier failure functions. The “failure” is the event of leaving unemployment thanks to a self-found job, so the Kaplan-Meier estimates censor observations when a provider-sourced job is found and at the last date of data extraction from the Department databases (February 5, 2019). We impose the usual assumption that censoring is ignorable, i.e., that time until exit from unemployment thanks to a self-found job (which is not fully observable) is independent of the censoring indicators – possibly conditionally on covariates. The first date at which “failure” can be observed is May 15, 2017, i.e., the day MJG went online and job seekers in the treatment group were directed to it. Given the randomized design, the failure function for the control group can be interpreted as the counterfactual job-finding rate that treated job seekers would have experienced in the absence of MJG at any point in time. The estimated failure (or hazard of leaving unemployment) functions are reported in Figure 4.

¹²The Australian Bureau of Statistics reports that in 2017, 98% of individuals aged 18-34 used the Internet regularly, followed by 94% of those aged 35-54, while a much lower 69% for those aged 55 or more.

Figure 4: Estimated hazard of exit from unemployment by treatment status



Notes: The figure shows unconditional Kaplan-Meier failure functions by treatment status. The horizontal axis measures the number of days since entering the trial; the “failure” is the event of leaving the unemployment state thanks to a self-found job. Sample: 2,690 workers (964 in age range 35–50; 1,045 of age 34 and younger; 681 of age 51 and older; 1,394 men; 1,296 women; 480 men in age range 35–50; 484 women in age range 35–50) who were unemployed and registered at the provider’s job centers at the start of the trial.

The figure shows that about 5% of all job seekers in the control group (who benefited from all the services of the provider except for referral to MJG) are matched within one month, about 10% within two months, and about 13% within three months. After six months, about 22% have obtained at least one job match, about 29% after one year, and about 36% after twenty-two months. A small positive difference between the job-finding rates of treated and control job seekers emerges after the second month. Such a difference is appreciably larger among job seekers between 35 and 50 years of age (possibly negative in the early months since assignment to treatment for older job seekers and after fifteen months for younger ones), especially women. For these women in the age range 35–50, the gap between the two failure functions is comparable to the gender gap in the job-finding rate that we observe after two months at the baseline, as discussed above in relation to the gender gap reversal. Table 2 reports the p -values from tests of the null hypothesis that the failure functions are equal, using the Log-rank and Wilcoxon tests. The results indicate that, over the entire duration of the evaluation period, the differences are statistically significant for job seekers in the age group 35–50 and particularly for women in this age group, in line with the results of the regression analysis.

Table 2: Test of the hypothesis that the hazard functions by treatment status are equal

Sample	Log-rank	Wilcoxon	N control	N treated
All job seekers	0.415	0.343	1,248	1,442
Job seekers age 35–50	0.070	0.044	451	513
Job seekers age 34 and younger	0.787	0.938	479	566
Job seekers age 51 and older	0.840	0.740	318	363
Male job seekers	0.848	0.840	628	766
Female job seekers	0.399	0.310	620	676
Male job seekers age 35–50	0.485	0.435	229	251
Female job seekers age 35–50	0.058	0.036	222	262

Notes: The table reports the p -values from tests of the null hypothesis that the Kaplan-Meier failure functions illustrated in Figure 4 are equal across treatment and control groups. The tests used are the Log-rank and Wilcoxon tests.

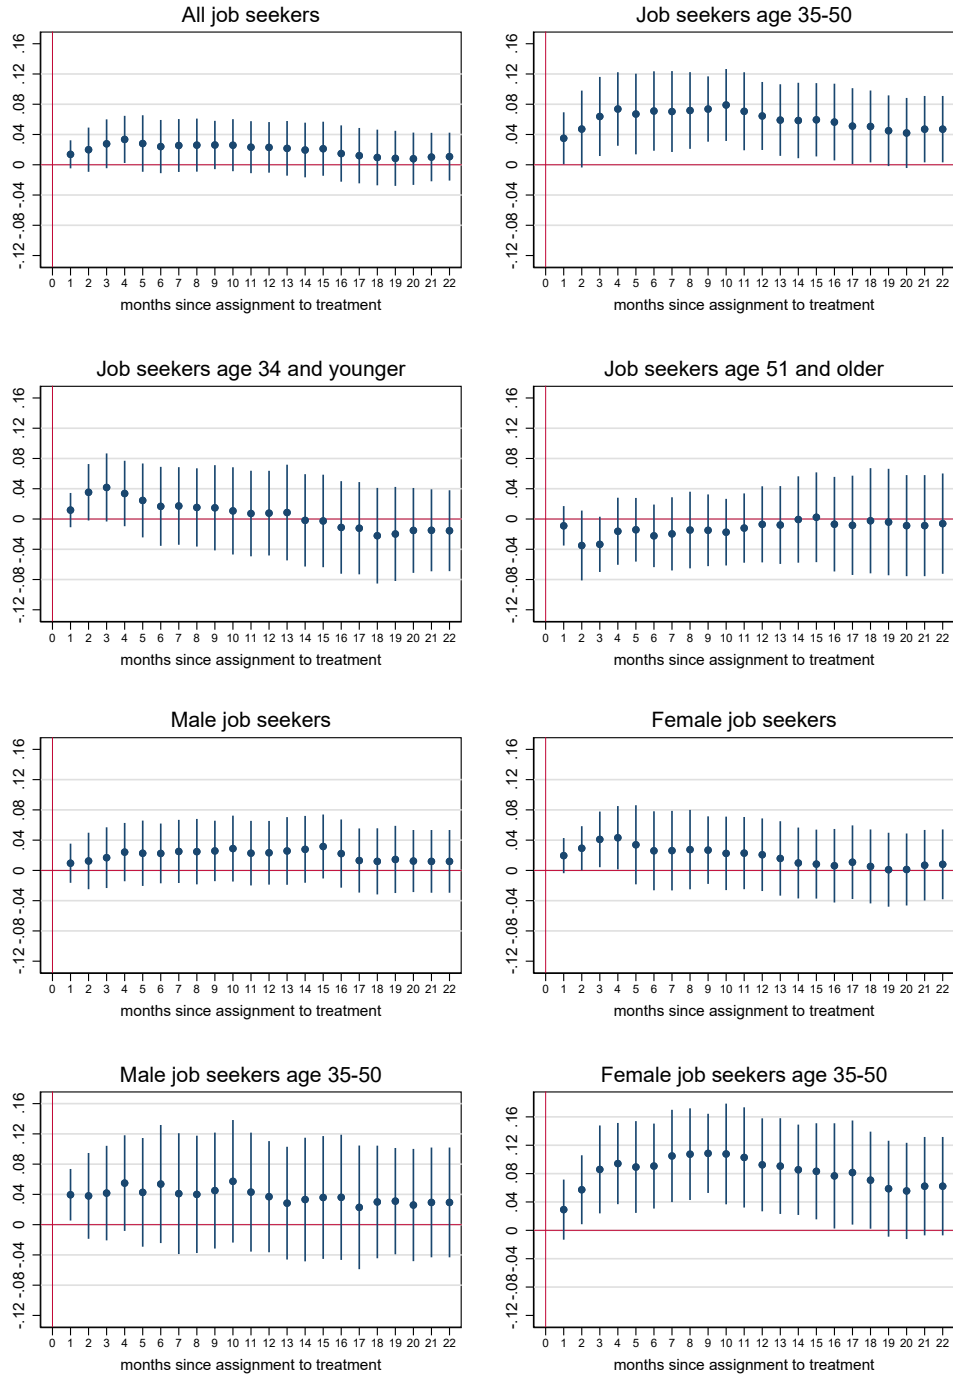
Next, we estimate a complementary log-log model so to provide average ITT estimates from a hazard model that can be more directly compared with those produced by OLS es-

timation of equation (1). The complementary log-log model allows us to estimate how the hazard of leaving the unemployment state is altered by assignment to treatment *each month*, conditional on being still unemployed during the previous month. The continuous-time hazard of exit from unemployment for job seeker i during month m is assumed to be $\tilde{\alpha}_m \exp(\tilde{\beta}_m D_i + \tilde{\gamma}_m X_i)$, where $\tilde{\alpha}_m$ is the baseline hazard – which we assume to be constant during month m – and a tilde is used to distinguish the parameters in equation (1) from those in this hazard function, which are distinct and yet are in direct correspondence. Thus, denoting by Y_i the month (since the beginning of the trial) during which a job seeker found a job, i 's probability of leaving unemployment during month m conditional on being unemployed during month $m - 1$ is given by

$$\Pr(Y_i = m | Y_i \geq m, D_i, X_i) = 1 - \exp[-\exp(\ln \tilde{\alpha}_m + \tilde{\beta}_m D_i + \tilde{\gamma}_m X_i)]. \quad (2)$$

The parameters of this equation are estimated, month by month, by ML starting from the first month of the trial, which is the first period when the hazard function is defined. The marginal effects (around the mean of the treatment indicator) associated with parameters $\tilde{\beta}_m$, for $m = 1, \dots, 22$, are the direct counterparts of the average ITT's reported in [Figure 3](#). Relative to the latter, the marginal effects from model (2) are more precisely estimated thanks to the parametric assumption that the duration of the unemployment state is exponentially distributed. [Figure 5](#) illustrates the estimates of these marginal effects and the associated 95% confidence intervals, which in this case are clustered at the job center level. This figure reproduces the post-treatment pattern of [Figure 3](#), with narrower confidence intervals. Such efficiency gain allows us to see more clearly the effects of MJG on job seekers in the age group 35–50 and women in this age range.

Figure 5: ITT effects on job finding, complementary log-log duration models



Notes: The figure shows, for each month since entering the trial the estimated average marginal Intention-To-Treat effect associated with parameter $\tilde{\beta}_m$ in equation (2), computed at the average value of the treatment indicator and represented by a circle. The outcome $\Pr(Y_i = m | Y_i \geq m, D_i, X_i)$ is the conditional probability that a job seeker leaves the unemployment state during month m conditional on being unemployed during month $m - 1$, thanks to *any* job that is classified by the job center as “Found Own Employment” (self-found). The figure also shows the 95% confidence interval associated with robust standard errors clustered at the job center level. Sample: 2,690 workers (964 in age range 35–50; 1,045 of age 34 and younger; 681 of age 51 and older; 1,394 men; 1,296 women; 480 men in age range 35–50; 484 women in age range 35–50) who were unemployed and registered at the provider’s job centers at the start of the trial.


3.3 Full-time and longer-lasting jobs

Next, we explore whether MJG had an effect on the type and duration of employment. We repeat the ML estimation of model (2) using more stringent notions of job finding. First, we consider matches with full-time jobs. Figure 6 shows that when pooling all job seekers (top-left panel), the average ITT effect on the transition rate from unemployment to a full-time job is positive and marginally significant after the first month from entering the trial (about 2 percentage points), increasing to a more precisely estimated 4 percentage points until the sixth month, before converging to an insignificant 2 percentage points. Among job seekers younger than 34 and among women in the age group 35–50, this effect is also statistically significant between the third month and the end of the first year, ranging between 3 and 6 percentage points.

Second, we consider job matches that, retrospectively, will last for at least 6 months.¹³ These more stable jobs are the hardest to secure for a job seeker at the baseline: after one year since entering the trial, only 15.6% of job seekers in the control group obtained a placement in such jobs, as opposed to 30% in any job. Figure 7 shows that in this case no significant ITT effects are detected. For some groups, the chances of securing a more stable job during the first months of the trial are so low that the ITT is a very precisely estimated zero. Therefore, MJG did not lead to a higher number of long-lasting jobs *relative to the control group*. Nonetheless, at the very least the intervention helps securing temporary job placements that may provide a springboard for future, more stable, employment (Booth et al., 2003; Ichino et al., 2008).

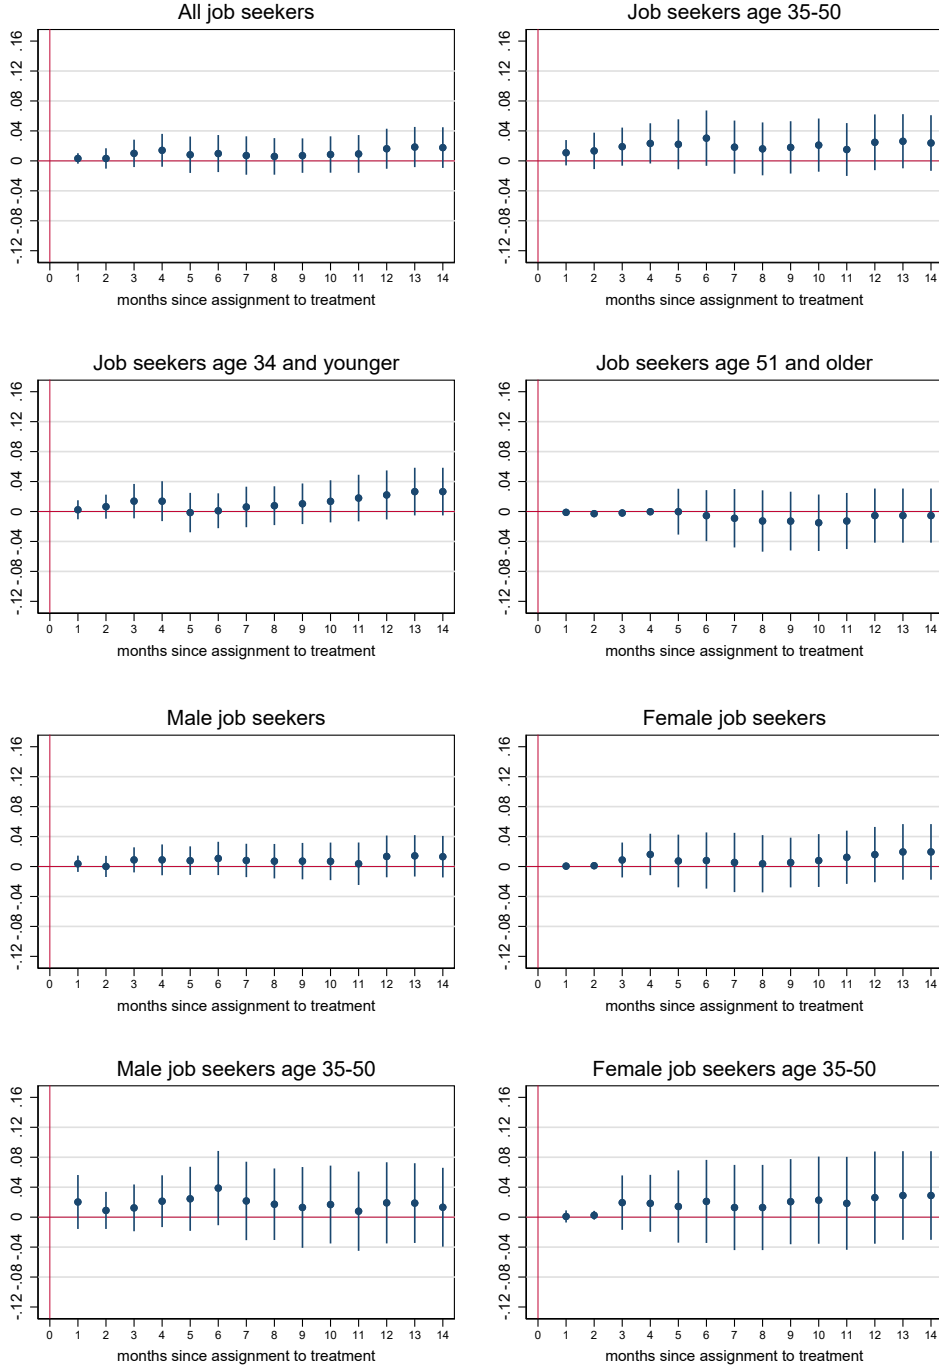
¹³Note that this outcome can only be observed for 14 months after the onset of the trial.

Figure 6: ITT effects on full-time self-found jobs



Notes: The figure shows, for each month since entering the trial the estimated average marginal Intention-To-Treat effect associated with parameter β_m in equation (2), computed at the average value of the treatment indicator and represented by a circle. The outcome $\Pr(Y_i = m | Y_i \geq m, D_i, X_i)$ is the conditional probability that a job seeker leaves the unemployment state during month m conditional on being unemployed during month $m - 1$, thanks to a *full time* job that is classified by the job center as “Found Own Employment” (self-found). The figure also shows the 95% confidence interval associated with the largest between the p -value from a test of $H_0 : \beta_m = 0$ produced by the wild cluster bootstrap (clustering at the job center level) and the p -value implied by heteroskedasticity-robust standard errors. Sample: 2,690 workers (964 in age range 35–50; 1,045 of age 34 and younger; 681 of age 51 and older; 1,394 men; 1,296 women; 480 men in age range 35–50; 484 women in age range 35–50) who were unemployed and registered at the provider’s job centers at the start of the trial.

Figure 7: ITT effects on self-found jobs kept for at least six months



Notes: The figure shows, for each month since entering the trial the estimated average marginal Intention-To-Treat effect associated with parameter β_m in equation (2), computed at the average value of the treatment indicator and represented by a circle. The outcome $\Pr(Y_i = m | Y_i \geq m, D_i, X_i)$ is the conditional probability that a job seeker leaves the unemployment state during month m conditional on being unemployed during month $m - 1$, thanks to a job that is classified by the job center as “Found Own Employment” (self-found) and that lasted *at least 6 months*. The figure also shows the 95% confidence interval associated with the largest between the p -value from a test of $H_0 : \beta_m = 0$ produced by the wild cluster bootstrap (clustering at the job center level) and the p -value implied by heteroskedasticity-robust standard errors. Sample: 2,690 workers (964 in age range 35–50; 1,045 of age 34 and younger; 681 of age 51 and older; 1,394 men; 1,296 women; 480 men in age range 35–50; 484 women in age range 35–50) who were unemployed and registered at the provider’s job centers at the start of the trial.

4 Conclusions

We evaluated whether a website that offers self-directed resources traditionally delivered by job centers' staff improves re-employment outcomes. The intervention increased the job-finding rates among individuals in the age range 35–50, particularly women. The gender difference in the ITT that we uncover reverts the gender gap in job-finding rates for this age group. We also observed a positive effect on full-time job placements driven by women and young job seekers. A simple cost-benefit analysis indicates that the benefits, in terms of reduced unemployment payments and increased income tax collection, range between about \$1,000 and \$1,350 per job seeker per month of anticipated exit from unemployment.¹⁴

The delivery of job search assistance in person is costly. Understanding whether components of this support can be provided online is therefore crucial for an efficient reallocation of resources towards more needy beneficiaries. While our experiment took place in the context of reemployment services, these share similarities with other publicly funded programs, which makes our findings more general. A move towards delivering welfare assistance online presents several advantages. First, both the cost of providing and receiving services would decrease. On top of the reasons that arose from the Covid-19 pandemic, this would be particularly so for time-constrained citizens, such as those with caring responsibilities, for whom turning up to regular mandated appointments can be costly and who would also prefer to access some of these services at their own pace. The online environment also offers a wide range of already freely available training platforms (e.g. MOOC or Coursera) that government services can help match to citizens based on needs. Second, such environment allows for program targeting at a lower marginal cost because policymakers can monitor take up and usage of online resources. Third, the online environment allows for a faster and more accurate evaluation of what works and for whom. Assisting beneficiaries online would enable policymakers to track intermediate outcomes, such as the level of engagement with training materials. This is a useful proxy to study job seekers' preferences and satisfaction with the assistance they receive, and it's available in real-time at virtually no cost. At the

¹⁴In 2017, the unemployment benefit job seekers could receive was between about \$850 and \$1,140 per month depending on their housing arrangement and family composition. The income tax collected from a minimum wage worker was around \$213 per month. Further details are available here <https://www.ato.gov.au/rates/individual-income-tax-rates/>

same time, our results suggest that not everyone may benefit equally from an online delivery mode. For these job seekers, support could continue to be provided face-to-face, even at a more intensive level thanks to the resources that are freed up by online delivery to other individuals. Such a hybrid support model for welfare beneficiaries – digital by default with access to in-person services when required – would be more tailored to users’ needs and preferences. For example, we found that some job seekers did not have an email address or phone listed, which could be a good proxy for redirecting them immediately to in-person services. Other potential indicators could include requesting an individual login by a certain deadline before being referred to offline services.

The usual caveats concerning randomized trials that test new policies apply to our experiment. It is possible that employment assistance interventions benefit treated job seekers at the expense of untreated ones, and so the beneficial impact that we detect may be smaller when general equilibrium effects are considered or if our small-scale intervention were scaled up (Crépon et al., 2013). Moreover, what works in Sydney may not necessarily work in other contexts. It should be noted though that our trial covers a geographical area that goes beyond the Sydney city center, including locations that are more than 40 miles from the economic hub of the city. Another special feature of our intervention is that it was implemented in Australia, a country with almost full high-speed Internet penetration. One could argue that in countries where geographical remoteness is associated with poorer broadband infrastructure the effect would be weaker. Moreover, a question that we cannot answer within our research design is whether online assistance is a complement or substitute for in-person services. Given that we could not deny job seekers the legislated compulsory in-person appointments, we cannot determine whether online and in-person assistance worked as complements or substitutes in producing the positive outcomes that we observed.

Despite these limitations, to the best of our knowledge this is the first study that provides experimental evidence from the field that providing online assistance to job seekers via resources that they can voluntarily use outside of job centers speeds-up reemployment. As such, the evidence that we have produced reinforces the case for online government assistance services as a promising avenue for both researchers and policymakers.

References

- Abadie, A., S. Athey, G. W. Imbens, and J. Wooldridge (2017). When should you adjust standard errors for clustering? Working Paper 24003, NBER.
- Altmann, S., A. Falk, S. Jager, and F. Zimmermann (2018). Learning about job search: A field experiment with job seekers in Germany. *Journal of Public Economics* 164, 33–49.
- Angrist, J., D. Lang, and P. Oreopoulos (2009). Incentives and services for college achievement: Evidence from a randomized trial. *American Economic Journal: Applied Economics* 1(1), 136–63.
- Angrist, J. D., G. W. Imbens, and D. B. Rubin (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association* 91(434), 444–455.
- Belot, M., P. Kircher, and P. Muller (2019). Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice. *The Review of Economic Studies* 86(4), 1411–1447.
- Booth, A., M. Francesconi, and J. Frank (2003). Temporary jobs: stepping stones or dead ends? *The Economic Journal* 112(480), 189–2013.
- Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics* 90(3), 414–427.
- Cameron, A. C. and D. L. Miller (2015). A practitioner’s guide to cluster-robust inference. *Journal of Human Resources* 50(2), 317–372.
- Card, D., J. Kluge, and A. Weber (2010). Active labour market policy evaluations: A meta-analysis. *The Economic Journal* 120(548), F452–F477.
- Card, D., J. Kluge, and A. Weber (2018). What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations. *Journal of the European Economic Association* 16(3), 894–931.
- Crépon, B., E. Duflo, M. Gurgand, R. Rathelot, and P. Zamora (2013). Do labor market policies have displacement effects? evidence from a clustered randomized experiment. *The Quarterly Journal of Economics* 128(2), 531–580.
- Crépon, B. and G. J. Van Den Berg (2016). Active labor market policies. *Annual Review of Economics* 8, 521–546.
- DJSB (2015). Employer feedback on lower skilled recruitment. Technical report, <https://docs.employment.gov.au/documents/employer-feedback-lower-skilled-recruitment>.
- Eckel, C. C. and P. J. Grossman (2008). Men, women and risk aversion: Experimental evidence. *Handbook of experimental economics results*, 1061–1073.
- Exley, C. L. and J. B. Kessler (2019). The gender gap in self-promotion. *National Bureau of Economic Research No. 26345*.

- Gürtzgen, N., A. Diegmann, L. Pohlen, and G. J. van den Berg (2021). Do digital information technologies help unemployed job seekers find a job? Evidence from the broadband internet expansion in Germany. *European Economic Review*, 1036–57.
- Hornsby, J. S. and B. N. Smith (1995). Resume content: What should be included and excluded. *SAM Advanced Management Journal* 60(1), 4–10.
- Ichino, A., F. Mealli, and T. Nannicini (2008). From temporary help jobs to permanent employment: What can we learn from matching estimators and their sensitivity? *Journal of Applied Econometrics* 23(3), 305–327.
- Kerry, S. M. and J. Martin Bland (2001). Unequal cluster sizes for trials in english and welsh general practice: implications for sample size calculations. *Statistics in medicine* 20(3), 377–390.
- Kuhn, P. and H. Mansour (2014). Is internet job search still ineffective? *The Economic Journal* 124(581), 1213–1233.
- Liebman, J. B. and E. F. Luttmer (2015). Would people behave differently if they better understood social security? Evidence from a field experiment. *American Economic Journal: Economic Policy* 7(1), 275–99.
- Morgan, K. L. and D. B. Rubin (2012). Rerandomization to improve covariate balance in experiments. *The Annals of Statistics* 40(2), 1263–1282.
- OECD (2012). *Activating Jobseekers. How Australia does it.*
- Romano, J. P. and M. Wolf (2005). Exact and approximate stepdown methods for multiple hypothesis testing. *Journal of the American Statistical Association* 100(469), 94–108.
- Rose, R. A. and G. L. Bowen (2009). Power analysis in social work intervention research: Designing cluster-randomized trials. *Social Work Research* 33(1), 43–52.
- Ross, C. M. and S. J. Young (2005). Resume preferences: Is it really “business as usual”? *Journal of Career Development* 81(396), 153–164.
- Rutterford, C., A. Copas, and S. Eldridge (2015). Methods for sample size determination in cluster randomized trials. *International journal of epidemiology* 44(3), 1051–1067.
- Schramm, R. and N. Dortch (1991). An analysis of effective resume content, format, and appearance based on college recruiter perceptions. *The Bulletin of the Association for Business Communication* 54(3), 18–23.
- Stevenson, B. (2009). The internet and job search. In *Studies of labor market intermediation*, pp. 67–86. University of Chicago Press.
- Thoms, P., R. McMasters, M. R. Roberts, and D. A. Dombkowski (1999). Resume characteristics as predictors of an invitation to interview. *Journal of Business and Psychology* 13(3), 339–356.
- Toth, C. (1993). Effect of résumé format on applicant selection for job interviews. *Applied HRM Research* 4(2), 115–125.