NBER WORKING PAPER SERIES

CAN A WEBSITE BRING UNEMPLOYMENT DOWN? EXPERIMENTAL EVIDENCE FROM FRANCE

Aïcha Ben Dhia
Bruno Crépon
Esther Mbih
Louise Paul-Delvaux
Bertille Picard
Vincent Pons

Working Paper 29914
http://www.nber.org/papers/w29914

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2022

We are grateful to Pierre-Louis Bithorel, Quiterie Landèche, Marion Richard, and Thomas Van Casteren for outstanding research assistance. We benefited from helpful comments and suggestions from Esther Duflo and Frank Schilbach. We thank Bayes Impact and Pôle emploi for collaborating on this experiment, and gratefully acknowledge funding from the "Fondation La France s'engage" ("Ministère de la Ville, de la Jeunesse et des Sports"). We received IRB approval from Harvard for this project (IRB17-0735). This work was supported by the French National Research Agency Grant ANR-17-EURE-0020, and by the Excellence Initiative of Aix-Marseille University - A*MIDEX. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w29914.ack

Can a Website Bring Unemployment Down? Experimental Evidence from France
Aïcha Ben Dhia, Bruno Crépon, Esther Mbih, Louise Paul-Delvaux, Bertille Picard, and Vincent Pons
NBER Working Paper No. 29914
April 2022
JEL No. D83,D84,J22,J24,J62,J64

ABSTRACT

We evaluate the impact of an online platform giving job seekers tips to improve their search and recommendations of new occupations and locations to target, based on their personal data and labor market data. Our experiment used an encouragement design and was conducted in collaboration with the French public employment agency. It includes 212,277 individuals. We find modest effects on search methods: the users of the platform adopt some of its tips and they are more likely to use resources provided by public employment services. However, following individual trajectories for 18 months after the intervention, we do not observe any impact on time spent looking for a job, search scope (occupational or geographical), or self-reported well-being. Most importantly, we do not find any effect on any employment outcome, whether in the short or medium run. We conclude that the enthusiasm around the potential for job-search assistance platforms to help reduce unemployment should be toned down.

Aïcha Ben Dhia
Massachusetts Institute of Technology
abendhia@mit.edu

Bruno Crépon
CREST
15, boulevard Gabriel Peri
92245 Malakoff Cedex
France
crepon@ensae.fr

Esther Mbih
CREST-Ensae
5 Avenue Henry Le Chatelier
Palaiseau 91764
France
mbih.esther@gmail.com

Louise Paul-Delvaux
Harvard University
lpauldelvaux@fas.harvard.edu

Bertille Picard
AMU - AMSE
5-9 Bd Maurice Bourdet
Marseille 13001
France
bertille.picard@univ-amu.fr

Vincent Pons
Harvard Business School
Morgan Hall 289
Soldiers Field
Boston, MA 02163
and NBER
vpons@hbs.edu

1. Introduction

Frictions and inefficiencies on the labor market account for an important share of overall unemployment. Even when jobs are available and job seekers and companies have good information, the matching process between them can take time and generate frictional unemployment (McCall (1970); Mortensen (1970); Michaillat (2012)). In addition, unemployed workers may face specific hurdles to access information on the demand side, process it, and translate it into an effective search strategy (Babcock et al. (2012); Cooper and Kuhn (2020)). Those who just started to look for a job may be overoptimistic about their chances and procrastinate (DellaVigna and Paserman (2005); Spinnewijn (2015); Mueller, Spinnewijn, and Topa (2021)). Those who have been looking unsuccessfully for a longer time should ideally widen the scope of their search and renew their efforts, but they may instead lose self-esteem and motivation, generating a vicious circle (Krueger and Mueller 2011).[1] Naturally, governments around the world invest considerable money to address these inefficiencies, and they are hungry for effective remedies.

In the late 1990s, hopes were high that the advent of the internet would lessen job market frictions by facilitating access to information (Autor 2009), but effects on employment long failed to materialize (Kuhn and Skuterud 2004; Kuhn and Mansour 2014). While online job boards increased unemployed workers' information on job openings, processing this growing amount of information also became more challenging. Therefore, a second set of online tools (sometimes incorporated in the same websites) advise job seekers on the types of jobs they should consider and, more broadly, on their search strategy. Several features make these tools appealing: they are relatively inexpensive, easy to scale, and can provide tailored recommendations informed by data on the labor market and on users' profiles. However, despite the increasing number and popularity of such tools, evidence on their effectiveness is limited.

In this paper, we evaluate the impact of "Bob Emploi," an online job-search assistance website established to reduce unemployment in France. Bob Emploi was launched in 2016 by Bayes Impact, a private non-profit organization. In designing the website, Bayes Impact benefited from a close collaboration with the French public employment agency, Pôle emploi, which records systematic information on each job seeker's periods of unemployment. Bob Emploi set out to analyze the rich data provided by Pôle emploi with state-of-the-art statistical analysis methods. This innovative partnership sets Bob Emploi apart from other private websites and increases its potential for impact. Bob Emploi endeavors to assist and motivate job seekers by providing personalized and data-informed 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. Our evaluation isolates the impact of these job-search assistance features since Bob Emploi does not include job listings.

Our experiment relates to a recent strand of the literature that tests narrower interventions, each focusing on a specific challenge faced during the job search. Belot, Kircher, and Muller (2019) evaluate an experimental

[1] Beyond informational and behavioral hurdles, important factors responsible for unemployment include occupational and geographical mismatch between vacancies and job seekers (e.g., Şahin et al. (2014); Marinescu and Rathelot (2018); Herz and Rens (2019)), as well as rigidities preventing wages from decreasing during recessions (e.g., Temin (1990); Campbell and Kamlani (1997)). It is beyond the scope of this paper to review this large literature in a systematic way.

platform helping job seekers expand their search. Using administrative data, the platform recommends alternative occupations that are close to subjects' existing skills and it provides information on the situation of the labor market in their regions of interest. This intervention broadens the set of jobs that unemployed workers consider and increases their number of job interviews. Altmann et al. (2018) send brochures motivating job seekers to search for employment by informing them about the consequences of unemployment. They report modest overall effects on employment and earnings, which are concentrated among individuals with a higher risk of long-term unemployment. Gee (2019) finds that communicating on the intensity of competition for jobs can also generate a change in behavior, namely an increased likelihood to complete an application. Finally, relying on the psychology literature, Abel et al. (2019) invite job seekers to complete a detailed job-search plan in order to help them reduce the gap between their intentions and behavior. They find that making a plan substantially increases job-search efficiency and employment.[2]

Our evaluation advances this literature in two important ways. First, Bob Emploi was not designed to focus on a single dimension of the job search, unlike the interventions mentioned above. Rather, it combines the different aspects investigated separately in these previous experiments, so our effects reflect possible complementarities between them. Second, Bob Emploi may not be as sophisticated on each dimension as previous interventions, but it is more representative of the content of websites used by large numbers of job seekers. Unlike the interventions examined in previous experiments, it was not designed for research purposes. While we generated exogenous variation in take-up, Bob Emploi is freely accessible across the entire country. To the best of our knowledge, we are the first to provide at-scale experimental evidence on the impact of an actual private website dedicated to job-search assistance. This is an important contribution, given the growing evidence that interventions which are effective in a small, controlled setting may become unimpactful when they are scaled up (e.g., Banerjee, Duflo, and Glennerster 2008; Banerjee et al. 2017; Grossman, Humphreys, and Sacramone-Lutz 2020).

To investigate whether using Bob Emploi changes job seekers' search strategies and, in turn, improves their re-employment prospects, we implemented a large-scale encouragement design. Our sample consists of 212,277 job seekers in 254 Pôle emploi local agencies. All job seekers included in the study had been registered with Pôle emploi for less than a year before inclusion in our sample. Randomization was conducted at the individual level, with 56.3% of the sample assigned to the treatment group. Individuals in this group were invited to attend an informational session at their local agency between April and July 2017, where Pôle emploi's caseworkers introduced them to Bob Emploi. In addition, all job seekers in the treatment group received three emails encouraging them to create a Bob Emploi account, whether or not they had attended the informational session.

Our measure of take-up of the treatment is a dummy equal to 1 if the job seeker created a Bob Emploi account or attended the dedicated informational session conducted by Pôle emploi, and 0 otherwise. Take-up reached 27.2% in the treatment group, and only 0.2% in the control group. While the information meetings focused on Bob Emploi,

[2] Beyond studies on the supply side of the labor market, our paper also relates to experiments studying the *demand side* (e.g., Horton 2017).

they also provided some information on other resources at Pôle emploi. However, the differential rate of usage between the treatment and the control groups was much larger for Bob Emploi than for other online resources mentioned at the meetings. Therefore, the effects of the intervention can mostly be attributed to using Bob Emploi.

We follow the trajectories of individuals in our sample for 18 months after the intervention. We rely on administrative records from Pôle emploi to measure socio-demographic variables, the rate of reemployment, and the use of Pôle emploi's support services by job seekers (e.g., meeting with a caseworker or attending a job training program). The data also record all job applications made through Pôle emploi's job board (also studied by Marinescu and Skandalis (2021)). In addition, we conducted an online survey six months after the end of the intervention to ask job seekers about the scope of their search (in terms of occupation and geographic location), the websites they used, and the time they spent searching every week. The survey also contained questions on respondents' well-being, whether they felt support or loneliness in their search, and activities outside of the job search. Even though the survey was conducted online, we obtained responses from 47,327 individuals (22.3% of the sample).

Overall, Bob Emploi had limited effects on job seekers' search strategy and no effect on reemployment outcomes.

First, we consider search efforts, and do not find any effect on the number of hours spent looking for a job every week. Analyzing the characteristics of job applications, we also see no change in job seekers' search scope, whether in terms of occupation type or geographic location. Survey respondents in the treatment group are not more likely to report sending unsolicited applications, and the effect on the total number of applications sent through Pôle emploi's job board is close to null and it is not significant. Job seekers may invest insufficient time in search efforts if they overestimate their chances to find a job. Comparing the time by which job seekers actually get a job with their original expectations, we find that 69.3% of job seekers in the control group were overoptimistic, echoing recent results, for example in Mueller, Spinnewijn, and Topa (2021). Despite the information provided by Bob Emploi, including a diagnosis of employability shared with each user, the website does not decrease this number.

Second, Bob Emploi does affect some of the methods used to search for jobs. Job seekers in the treatment group are slightly more likely to report using some of the best practices frequently recommended by Bob Emploi (using their personal and professional network and following up with recruiting firms after sending their application), and the number of Pôle emploi websites they use to look for jobs increases by 4.7%. By contrast, the number of private websites used (beyond Bob Emploi) remains unchanged. Moreover, individuals in the treatment group are 2.4 percentage points more likely to meet with their Pôle emploi caseworker in the six months following the intervention, and the website did not decrease participation in other search assistance programs, such as group workshops or training programs. These results suggest that Bob Emploi acts as a complement, rather than a substitute, to public job-search assistance services.

Third, to provide a more complete picture of the potential mechanisms through which Bob Emploi may affect job-search effectiveness, we turn to survey variables on well-being and life balance. Periods of unemployment can be stressful, especially when job seekers feel isolated, poorly advised, or poorly equipped to find work. The failure to get called back for interviews and find a job can generate important psychological costs, drain the job seeker of energy for social activities, and demotivate their future search efforts (e.g., McKee-Ryan et al. 2005; Krueger and Mueller 2011; Krueger and Mueller 2012). Bob Emploi attempts to assist job seekers in overcoming these psychological barriers by engaging them with a user-friendly interface, sending them uplifting messages, and encouraging them to engage in activities outside the job search that might brighten their outlook. Survey respondents in the treatment group are somewhat more likely to report feeling supported in their search for a job. Surprisingly, however, they are slightly *less* likely to say that they are motivated in their search, compared to job seekers in the control group. They do not report higher well-being and are no more likely to engage in sport, cultural, or volunteering activities.

Fourth, and most importantly, we investigate whether the modest effects on job search translated into effects on employment. As we observe all exits and entries in and out of unemployment, we know whether individuals were employed at any point following the intervention and we can compute the total number of days spent unemployed. We do not find any impact of Bob Emploi on any employment outcome, at any time horizon. Bob Emploi did not increase the likelihood of experiencing any employment episode, whether stable or not, and it did not reduce the total duration of unemployment. Due to the large scale of our experiment, our estimates are precise nulls. For instance, considering the upper bound of the 95 percent confidence interval, we can reject any effect higher than 0.4 percentage points on experiencing some employment episode within the 18 months following the intervention and any reduction of the time spent unemployed larger than two days. To shed light on the cost-benefit analysis of this type of job assistance platform, we also consider the amount of unemployment benefits received by job seekers. Again, we do not observe any significant difference between the treatment and control groups.

Finally, we conduct heterogeneity analyses to rule out the possibility that our zero mean effects on employment hide positive and negative effects for different groups of individuals. We first consider heterogeneous effects by age and education by explicitly interacting the treatment dummy with these two characteristics. We then explore heterogeneity across a larger set of characteristics by using the generic machine learning approach developed by Chernozhukov et al. (2018). We do not find any heterogeneous effect on unemployment duration six or 18 months after the intervention using either method.

In sum, Bob Emploi did not fulfill its promise of a breakthrough in job-search assistance, in spite of its close partnership with the French public employment services and access to their rich data. Beyond the performance of this particular website, our results raise questions about the potential for automated online platforms in general to reduce unemployment by advising job seekers based on large-scale data analysis. The sobering null effects we find across the board on employment contrast with the positive effects found on job interviews and offers, at least for some subgroups, by Belot, Kircher, and Muller (2019), Altmann et al. (2018), and Abel et al. (2019). One possible explanation is that researcher involvement in the design of these interventions enabled a higher quality of program content. While less impactful, the types of advice given by Bob Emploi are more representative of the online assistance commonly available to job seekers. Our results indicate that online resources of this type may affect job-search methods but that they are insufficient to significantly reduce unemployment.

2. Research design

2.1. Online resources for job seekers in France

The French Public Employment Agency, Pôle emploi, is responsible for distributing unemployment benefits to unemployed workers, assisting them and tracking their job searches, and helping match them with employers. With the rise of the internet at the end of the 1990s, private websites emerged as alternative labor market intermediaries. The first generation of these comprised simple job boards such as CareerBuilder or Monster (Kuhn and Skuterud 2004). A second wave leveraged the power of social networks, including websites with international coverage such as LinkedIn and Viadeo, as well as national (e.g., RegionsJob in France) and sector-specific websites (e.g., emploi-environnement.com for the environmental sector in France). At the same time, Pôle emploi started to digitize its services and created its own online platform, called "Emploi Store." This platform proposes online services from Pôle emploi or the Ministry of Labor, such as "La bonne formation" which helps users find training programs matching their needs and "#1jeune 1 solution: la boussole" which guides them in identifying benefits they are entitled to. In addition, Emploi Store publishes a list of private websites deemed useful for job seekers.

Beyond listing job openings, private websites often provide some form of job-search assistance, from freely accessible documents listing best practices when searching for a job[3] to fee-based coaching services. The advent of artificial intelligence and machine learning techniques over the last decade has opened up new opportunities to leverage large datasets and provide customized recommendations to job seekers and recruiters. Many websites charge employers for advertising their job listings to selected job seekers and for giving them automated assistance in sorting applications, using algorithms trained on data from their past hirings and information scraped from applicants' résumés. Fewer solutions have been designed to optimize the job search from the perspective of the job seeker. As in other fields of application, big data algorithms are only as good as their number of users and the quality of their data. But obtaining systematic data on job search and employment outcomes can represent a significant challenge for private websites. Bayes Impact occupies a unique position from that perspective, since it has access to comprehensive datasets on job seekers' trajectories through Pôle emploi.

2.2. Bayes Impact and Bob Emploi

Bayes Impact is a non-profit organization created in 2014 with the goal of leveraging big data analysis techniques to help resolve important public policy challenges and modernize public services. Based in Paris and San Francisco, the organization is financed almost entirely by public funding and builds partnerships with public institutions in order to access relevant data. For example, in 2017 it launched a project to document the use of force by the police, in partnership with the police services in California.

[3]For example, RegionsJob in France provides recommendations in dedicated posts such as "how to write a good résumé?" or "8 recommendations to perform well in an interview."

In France, Bayes Impact formed a partnership with Pôle emploi in 2016 – a time when nationwide unemployment had reached 10% – to obtain data on the job-search trajectories of registered job seekers and hiring conditions in the job market. In November 2016, Bayes Impact launched the online platform Bob Emploi, which combines Pôle emploi data with information provided by job seekers and recommends customized actions with the aim of increasing their chances of finding a job. Bayes Impact's C.E.O. set high expectations by announcing that the website could lower the unemployment rate by up to 10%.[4] Since its creation, Bob Emploi has received much attention from the media and continuous support from policymakers, including in the form of public aid.[5]

Bob Emploi extracts job market information from three databases from Pôle emploi.[6] The first of these is "Fichier historique," in which Pôle emploi records data on all periods of unemployment over the previous ten years for all people ever enrolled, in order to be able to determine the unemployment benefits they are entitled to and monitor their job search. The "Fichier historique" also contains socio-demographic characteristics of the job seekers (their gender, age, level of education, and municipality of residency), the job-search strategy they adopt at the beginning of their unemployment (their target sector, reservation wage, and skill level of the desired occupation),[7] and the Pôle emploi activities in which they participate (e.g., meetings with caseworkers and job training programs). Bayes Impact obtained access to a representative sample of the "Fichier historique." While this dataset keeps comprehensive data on periods of unemployment, it does not contain any information about the type of job (occupation, location, and type of contract) obtained by previously unemployed workers. The second database contains job offers posted by employers on Pôle emploi's job board. Each observation in the database indicates the job title and description as well as the number of applications received. This information provides a proxy for market tightness for each occupation in each location. It also allows Bayes Impact to identify the skill requirements of different types of jobs, using text analysis. The third dataset is an annual survey of a sample of 500,000 employers conducted by Pôle emploi to learn about their hiring needs in the six subsequent months and the skills they deem particularly difficult to recruit ("Besoins en Main d'Oeuvre"). Pôle emploi uses this survey to measure the labor demand across sectors and orient job seekers towards high-demand sectors.

Using these rich data, Bob Emploi seeks to help job seekers find a position through four different means. These four features target important aspects of job search behavior (see Cooper and Kuhn 2020 for a recent review of the literature on behavioral job search) and they echo previous interventions tested in the field (Belot, Kircher, and Muller 2019; Altmann et al. 2018; Abel et al. 2019).

[4] See e.g., *Europe 1*, "Paul Duan : "Je pense qu'on peut baisser le chômage de 10%"", https://www.europe1.fr/societe/paul-duan-je-pense-quon-peut-baisser-le-chomage-de-10-2901731, accessed July 2021.

[5] See e.g., *Les Echos Start*, "Bob Emploi, l'algorithme qui s'attaque au chômage", https://start.lesechos.fr/innovations-startups/tech-futur/bob-emploi-lalgorithme-qui-sattaque-au-chomage-1176667, accessed July 2021 and *Europe 1*, "Cinq choses à savoir sur Bob Emploi, le site qui veut enrayer le chômage", https://www.europe1.fr/economie/cinq-choses-a-savoir-sur-bob-emploi-le-site-qui-veut-enrayer-le-chomage-2901977, accessed July 2021.

[6] Additional details about these databases are published by Bayes impact on GitHub: https://github.com/bayesimpact/bob-emploi/blob/master/data_analysis/data/README.md#bmo-bmo-besoin-en-main-doeuvre, accessed July 2021.

[7] The sector is recorded as a five-digit occupation code based on Pôle emploi's lexicographical classification of occupations (called Répertoire ROME), which is an equivalent of the US O*NET classification. A five-digit occupation corresponds to a narrow type of job, such as nurse-anesthetist or hotel receptionist.

First, Bob Emploi gives personalized advice based on user profile (see examples in Appendix Figures A1 and A2). When creating a profile on Bob Emploi, each user is asked to fill in their socio-demographic characteristics (gender, age, level of education, and municipality of residency), their desired occupation, and their target geographic area and salary. Based on this information, and leveraging the labor market information provided by Pôle emploi, the website presents the user with a diagnosis of employability in the form of a numerical score based on the labor demand for the target job, the competition for these jobs, and the quality of the match between the user's profile and the job requirements. The idea underlying this feature, which is supported by a growing literature, is that job seekers tend to have inaccurate beliefs about the situation of the labor market and to be overoptimistic about their return to employment (Mueller, Spinnewijn, and Topa 2021), leading to suboptimal search strategy. The diagnosis of employability was expected to lead job seekers to update their beliefs about their chances to find a job and to adjust their search accordingly. The motivation for the intervention informing job seekers about the consequences of unemployment in Altmann et al. (2018) is similar.

In addition to the employability diagnosis, the website gives users recommendations to undertake some actions, ranked by order of importance. For instance, Bob Emploi's algorithm may advise users targeting a tight market to explore related careers with higher demand. If the data indicate that the market for that occupation is less tight in neighboring municipalities, users may be advised to widen the geographic perimeter of their search. Similarly, Belot, Kircher, and Muller (2019) encourage job seekers to broaden the set of jobs they consider. An important difference is that, unlike Bob Emploi, their tool exposes users to actual job openings.

To customize its recommendations, Bob Emploi also asks users to indicate how many job opportunities they identified over the past weeks, how many applications they sent, and how many interviews and job offers they received. The website uses the ratio between any two of these indicators to customize pieces of advice focusing on steps to overcome the most important obstacles to finding a job. For example, if the number of interviews received by a user is disproportionately low relative to their number of applications, Bob Emploi may advise them to edit their résumé or to improve other parts of their application file.[8]

Second, Bob Emploi provides general advice, such as how to behave in a job interview (Appendix Figure A3). Recommendations are frequently completed by referral to external resources, most often also online. Both the general and personalized advice provided by the website may help job seekers overcome informational barriers and increase their chances to find a job.

Third, beyond information provision, Bob Emploi also provides step-by-step planning assistance to its users (Appendix Figure A4). The goal of these action plans is to help job seekers organize their job search efficiently and prioritize the tasks they need to complete. Encouraging individuals to follow such plans has proven effective

[8]The algorithms used by Bob Emploi to provide tailored advice to job seekers are based on recommendations that Bayes Impact gathered through interviews with professional recruiters, Pôle emploi caseworkers, and career coaches. The initial objective of Bayes Impact was to identify the most important determinants of job-search success using machine learning techniques, but the data provided by Pôle emploi were insufficient to achieve satisfactory predictive power, due to the lack of specific information about the type of job obtained by former unemployed workers, in the "Fichier historique," and to the absence of systematic data on the applications sent and the search methods used by job seekers present in that dataset.

in many contexts, including job search: Abel et al. (2019) find that completing a detailed job search plan increases the number of job applications submitted and the likelihood to find a job.

Fourth, Bob Emploi sends regular reminders and messages of encouragement to job seekers as a motivation technique. Users can control how often they receive these emails. This feature is motivated by recent research in behavioral economics showing that present biased time preferences and biased beliefs can otherwise lead job seekers to procrastinate their search effort (DellaVigna and Paserman (2005); Spinnewijn (2015)). In addition, the website also generally uses encouraging language and an optimistic tone, emphasizing the strengths of each job seeker's profile. The format of the interface is simple and inviting, in contrast with the often austere and complex appearances of other administrative resources. These features may be particularly relevant given the decline in search effort and motivation over the unemployment spell and the increase in sadness documented by Krueger and Mueller (2011).

The four functionalities of Bob Emploi are representative of features proposed by other websites providing job-search assistance, as shown in Table 1. This table compares Bob Emploi with four other French websites as well as three U.S. websites. Some of these provide job listings and matching services, while others, like Bob Emploi, do not. But all websites shown in the table include at least one of the job-search advice or assistance features offered by Bob Emploi.

All four French websites used for comparison are listed on Pôle emploi's Emploi Store, like Bob Emploi. *Meteojob* (3,000,000 registered users)[9] provides general advice about CV writing and a list of training programs proposed by different organizations. Other websites go one step further and give tailored job-search advice, similarly to Bob Emploi. For instance, *Jobijoba* (2,500,000 users per month)[10] uses personal information from users' CVs to identify jobs suited to their profile as well as the wage they can aim for. Similarly, *Monkey tie* (500,000 users)[11] identifies jobs that users might be interested in based on their skillset and personality. Finally, other websites exclusively offer planning assistance to keep job seekers motivated. For instance, *MEMO* (115,000 users)[12] supports job seekers throughout the application process. Because job seekers often send applications to many companies without following up, *MEMO* classifies the applications and sends emails reminding job seekers where they stand in the application process and when they should call employers back.

In the U.S., the website *Indeed* (250 million visitors every month)[13] mainly focuses on matching services by displaying job offers people can apply to.[14] However, it also provides general information such as the names of

[9] https://www.meteojob.com/qui-sommes-nous, accessed July 2021.

[10] *Challenges*, "Cette start-up propulse l'emploi dans l'ère du big data", https://www.challenges.fr/entreprise/cette-start-up-propulse-l-emploi-dans-l-ere-du-big-data_119281, accessed July 2021.

[11] https://www.monkey-tie.com/, accessed July 2021.

[12] https://www.pole-emploi.fr/candidat/vos-services-en-ligne/emploi-store/memo--le-tableau-de-bord-intelli.html, accessed July 2021.

[13] https://www.indeed.com/about, accessed July 2021.

[14] U.S. websites compared to Bob Emploi were identified by searching for U.S. apps on the AppStore. We searched for the following words: "job," "job offers," "job search," "career development," "career coaching," "planning assistance," "professional development," "career advice," and "job seeker help." For each of these words, we selected the first 50 apps on the results page of the AppStore. We then dropped games and apps which had fewer than 100 user ratings, we restricted the sample to apps listing "Business" or "Productivity" as a characteristic, and we kept the apps with the most ratings. We analyzed the resulting list of 41 apps by hand, to identify apps

firms that are recruiting and the wages to expect from different types of jobs, and it relies on motivational content. For example, videos posted on its Youtube channel state that "it is 100% possible" for its users to find their dream job. The website *ZipRecruiter* (110 million users)[15] also includes a matching service, but goes one step further by providing personalized advice, such as the types of job users might be interested in based on the content of their résumé and their behavior on prior visits on the website. Finally, other websites encompass planning assistance. For instance, *Snagajob* (100 million users in the U.S.)[16] lists applications the user has made on the website and it marks completed applications, in-progress applications (for which the website sends reminders to complete the process), and applications for offers that expired.

2.3. Sampling frame

Together with Pôle emploi, we selected a total of 254 local agencies to participate in the experiment. These agencies were drawn randomly within strata, to ensure representativeness in terms of size and geographic location.[17]

Within these agencies, job seekers included in the sample were selected based on five eligibility criteria: (i) having been unemployed for a year or less at the moment of the sampling;[18] (ii) having declared oneself available to start working immediately;[19] (iii) having identified a target occupation and location and demonstrated autonomy in one's job search;[20] (iv) not having worked part-time hours exceeding half of full-time employment per month; and (v) having a valid email address and agreeing to receive informational emails from Pôle emploi. These criteria were defined by a working group of Pôle emploi caseworkers involved in the design of the experiment, to target people most likely to be interested in Bob Emploi and benefit from it.

The intervention was rolled out over a three-month period, from April to June 2017. On the first day of each month, up to 320 job seekers were drawn in each local agency from up-to-date lists, and randomly assigned to the treatment and control groups. In some (smaller) agencies and months, the number of eligible job seekers was lower than 320. In that case, 160 of the eligible job seekers were randomly allocated to the treatment group and the remaining ones to the control group. For instance, with a total of 293 eligible job seekers in an agency x month, 160 would be allocated to the treatment group and 133 to the control group. Overall, our final sample includes 212,277 individuals: 119,525 in the treatment group and 92,752 in the control group. All our specifications include agency x month fixed effects to account for small variations in the likelihood to be treated across agencies and over time.

providing job-search assistance. We visited all corresponding websites and selected those most relevant for the comparison with Bob Emploi.

[15] https://www.ziprecruiter.com/about, accessed July 2021.

[16] https://www.snagajob.com/about/, accessed July 2021.

[17] Furthermore, we only selected local agencies with the capacity to organize weekly informational sessions throughout the three months of the experiment.

[18] Unemployment duration was computed starting with job seekers' last registration at Pôle emploi. It does not take prior unemployment periods into account.

[19] This criterion served to exclude people enrolled in a training program or an internship and those not actively looking for a job due to health reasons.

[20] When job seekers register at Pôle emploi, caseworkers assess their level of autonomy based, for instance, on their ability to browse the internet with ease, and they assign them to one of three assistance tracks. Job seekers deemed autonomous in their job search receive more distant monitoring and assistance from caseworkers.

2.4. Encouragement design

Our experiment relied on an encouragement design. Job seekers in the treatment group were invited to an informational session at their local agency. Such sessions are frequently organized by Pôle emploi caseworkers to present free tools and programs to job seekers. In the present case, each meeting was hosted by one or two Pôle emploi caseworkers and lasted approximately 90 minutes. It began by introducing the challenges associated with a job search and the value of widening one's search methods, notably through personal and professional networking. The caseworkers then briefly presented the "Emploi Store," the Pôle emploi platform that lists job-search services, including Bob Emploi. The session then turned to the presentation of Bob Emploi using screenshots and, sometimes, a real-time simulation. Participating local agencies organized these information meetings at a rate of approximately one per week from April 20 to July 31, 2017. Job seekers in the treatment group were invited to these sessions based on the month in which they were drawn. Specifically, individuals from the first wave, drawn on April 1, were invited to attend meetings taking place between late April and late May. Individuals from the second wave, drawn on May 1, and third wave, drawn on June 1, were invited to attend meetings taking place between late May and late June, and between late June and late July, respectively.

To maximize take-up, we completed the informational sessions by sending three emails encouraging all job seekers in the treatment group to create an account on Bob Emploi. The emails were sent on July 7 and 26 (for individuals drawn in the first two months and in the third month, respectively), September 28, and November 13. (See Figure 1 for the full timeline of the experiment.) The exact text of the emails differed slightly for recipients who had attended an informational session and those who had not.

3. Data

To measure outcomes, we collected administrative data from Pôle emploi and administered an online survey. In this section, we provide detailed information on the outcomes constructed using these two data sources and show descriptive statistics and balance checks.

3.1. Administrative data from Pôle emploi

The administrative data from Pôle emploi provide detailed information about job seekers from their registration with the agency to the end of their unemployment period.

Unemployment records

We first use Pôle's emploi unemployment records (called "Fichier historique"), which are also used by Bob Emploi and were thus already described in Section 2.2. The individual monthly unemployment benefits reported in these data allow us to compute the total benefits that job seekers received over the six or 18 months following the intervention.

Variables describing an individual's target job are entered only at the beginning of the period of unemployment, so they cannot be used to detect changes resulting from the intervention.

Employment outcomes

We identify employment episodes based on the ICT01 employment indicator of Pôle emploi, which draws on employer declarations at the time of hiring and on job seekers' self-reported declaration of joining the labor force. That second source of information is necessary to track formerly unemployed workers who have become self-employed. As described in Section 2.3, we drew our sample in three separate installments. Accordingly, we compute employment outcomes starting from the relevant drawing date: April 1, May 1, and June 1 for the first, second, and third waves, respectively.

Our first employment outcomes measure whether job seekers have found a job six and 18 months after the intervention. Since informational sessions took place within two months after the drawing date (e.g., until the end of May, for individuals drawn on April 1), we define employment six months after the end of the intervention as a dummy equal to one for job seekers with any employment episode in the eight months following the drawing date (e.g., between April 1 and December 1 for individuals in the first wave). Employment 18 months after the end of the intervention is defined similarly. We consider two additional outcomes to characterize the length and stability of employment episodes. First, we count the total number of days in unemployment since the sampling date. Second, people can be employed under a long-term contract (called "Contrat de Durée Indéterminée," or CDI), a short-term contract (called "Contrat de Durée Déterminée," or CDD), and for part-time work. We measure effects on employment in stable jobs – that is, whether individuals obtained any work contract of more than six months (CDI or CDD longer than six months) over the period of interest.

Use of Pôle emploi's assistance services

Pôle emploi records the following actions completed by job seekers: one-on-one meetings with their caseworkers, as well as participation in workshops (e.g., résumé-writing classes), in individual assistance programs (e.g., skill assessment programs), or in any informational session such as the ones organized for the purpose of this experiment. Based on these records, we create two variables measuring respectively the number of meetings with the Pôle emploi caseworker and the total number of workshops, programs, and meetings that the job seeker attends following the intervention. Estimating effects on these outcomes enables us to investigate potential substitution effects between Bob Emploi and Pôle emploi.

Online applications on Pôle emploi's job board

We use individual-level data about applications sent through the job board administered by Pôle emploi. This job board is one of the most popular search platforms in France, covering about one third of all job creations. While vacancies on the website are mostly low-skilled positions, Skandalis (2019) finds they are more often for permanent contracts and with better pay than average jobs created in France.

Since 2014, job seekers can apply to these job listings online, which gives us access to micro-level data on job applications. We first count the total number of applications made on this job board and break them down into three sub-categories, depending on whether the application was initiated by the job seeker, their caseworker, or the employer.

The data also include some of the job listing features such as the job title and its geographical location. Based on this information, we create variables capturing the geographical scope and the occupational scope of the job search: the number of applications outside the job seeker's municipality, and the number of applications outside their target sector.

3.2. Survey data

To collect additional information on job-search behavior, we administered an online survey to the entire sample on 23 January 2018, six months after the end of the intervention. The full questionnaire is available in Appendix A.2. We first recorded information on job seekers' search effort, as measured by the number of hours they spent searching in a normal week and the frequency at which they submitted unsolicited applications. Second, we inquired about their job-search methods, namely whether they relied on their personal network and on other types of contacts;[21] how many search websites they used (including those created by Pôle emploi and private websites);[22] whether they looked for jobs in their municipality, département, or region of residency; and whether they followed some of the best practices frequently recommended by Bob Emploi to increase search efficiency (adapting one's CV or cover letter by reusing words or expressions of the job listing, and following up with recruiting firms after sending one's application). Third, we asked respondents to indicate their general well-being level, their motivation, and the extent to which they felt supported in their job search on 0 to 10 Cantrill scales. To gather additional information about their life balance, we also asked them if they participated in any sport, art, or community activity at least once a month. Finally, we asked them when they expected to find a job and used the answer to create a dummy variable for responses falling within the following three months.[23]

The first question was included in the email inviting people to respond to the survey, and the next questions were

[21] Specifically, we asked job seekers if they relied on friends and family, friends of friends, former classmates, former colleagues, their Pôle emploi caseworker, volunteers from organizations assisting job seekers, local business owners, and/or people found online.

[22] We provided a list of nine websites, including Bob Emploi. For each of them, we asked respondents whether they knew of the website, whether they used it, and whether they found it useful.

[23] All questions were adapted to the employment status of the respondent. For respondents who had found a job by the time of the survey, all questions focused on their latest job-search period, and questions about how much time they expected to need to find a job were removed.

administered through SurveyMonkey. As shown in Appendix Table A1, the fraction of job seekers who responded to at least one question beyond the first was 22.3% overall (for a total of 47,327 respondents), which compares favorably with other online surveys of job seekers.[24] We note a small (0.5 percentage points) but statistically significant difference between the control and treatment groups.[25]

3.3. Verifying randomization

Table 2 presents summary statistics for individuals in the sample. Column 1 shows the mean of pre-intervention variables for individuals in the control group and column 2 the results of regressions testing for differences in means with the treatment group. Only one out of 12 differences is significantly different from zero (at the 10% level), which confirms that the randomization was successful at balancing the treatment and control groups.

Most job seekers in the sample are aged 25 to 54 years old, and half of them are women. The requirement that job seekers be autonomous in their job search to be eligible for inclusion in the study resulted in only 16% of the sample having no high-school diploma, while 31% have a vocational degree, 19% an end-of-high-school diploma, and more than a third a university degree. Finally, recall that the sample was restricted to job seekers unemployed for a maximum of one year. More than half of them had been registered at Pôle emploi for more than six months.

Appendix Table A2 investigates the representativeness of the survey sample along the exact same dimensions. Survey respondents are more likely to be women than non-respondents and they are significantly older and more educated. They had also been unemployed for a slightly longer time by the time of the survey.[26]

4. Empirical specifications

4.1. Intent-to-Treat estimation

We measure the impact of having access to Bob Emploi on job-search strategy, employment outcomes, and well-being of unemployed individuals. For each outcome, we estimate two parameters of interest: the Intent-to-Treat (ITT) effect of being assigned to the treatment group and the Two-Stage Least Squares (2SLS) effect of actually opening a Bob Emploi account and being treated.

ITT effects are estimated using the following equation:

[24] For instance, the response rate to online surveys administered by Pôle emploi usually hovers around 10%, and DellaVigna et al. (2020) obtain a response rate of less than 8% in their survey of unemployed workers in Germany.

[25] The number of respondents varies across outcomes since people could refuse to respond to some of the questions or stop responding altogether. As shown in Appendix Table A1, column 2, the difference between the fraction of respondents who answered all unfiltered questions in the control and treatment groups is smaller than the difference in responding to any question beyond the first (column 1), but it remains significant.

[26] In Appendix Table A3, we check whether the effects on employment outcomes measured on the subsample of survey respondents are similar as the effects in the full sample. As in the full sample, all effects are nonsignificant, with only one exception: if anything, unemployment duration was *longer* among respondents in the treatment group. However, this effect is small (0.4% and 0.6% six months and 18 months after the intervention, respectively) and only significant at the 10% level. Furthermore, the effect on self-declared employment reported by survey respondents is small, negative, and nonsignificant, similarly as the effect measured in the same sample, using administrative data.

$$Y_{it} = \alpha_1 + \beta_1 T_i + \mu_{i1}^{a,m} + \epsilon_{it1}, \tag{1}$$

where Y_{it} is the outcome of interest for individual i at time t, T_i is a dummy equal to one for individuals in the treatment group, and $\mu_{i1}^{a,m}$ are strata fixed effects indicating the local agency a x month m of inclusion in the sample. As explained in Section 2.3, we include these fixed effects to account for small variations in the likelihood of being treated across agencies and over time. Conditional on these fixed effects, random treatment assignment ensures that T_i is orthogonal to ϵ_{it1}.

The key coefficient of interest, β_1, indicates the effect at time t of being assigned to the treatment group. We present results for t = six months and 18 months after the intervention for our employment outcomes. We also provide graphical evidence on changes in treatment effects on these outcomes over time. Throughout the analysis, we adjust standard errors for clustering by local agency.[27]

4.2. Definition of take-up and 2SLS estimation

2SLS effects are estimated using the following specification:

$$Y_{it} = \alpha_2 + \beta_2 B_i + \mu_{i2}^{a,m} + \epsilon_{it2}, \tag{2}$$

where B_i is a dummy equal to 1 if the individual was actually treated and is instrumented with T_i in the following first stage regression:

$$B_i = \gamma + \delta T_i + \mu_{i0}^{a,m} + \eta_i. \tag{3}$$

The coefficient β_2 indicates the effect at time t of being actually treated.

Importantly, informational sessions may have affected outcomes independently of opening an account on Bob Emploi since they gave job seekers an opportunity to interact with their caseworker and to learn about other resources provided or referenced by Pôle emploi on the "Emploi Store." Therefore, setting the take-up variable B_i to 1 only for individuals registered on Bob Emploi could lead to violations of the exclusion restriction. Instead, we also set B_i to 1 for individuals who participated in an informational session.

The list of individuals who created an account on Bob Emploi is obtained by matching email addresses of Bob Emploi users, provided by Bayes Impact, with email addresses of job seekers in the sample, provided by Pôle emploi. For privacy reasons, Bayes Impact could not directly share user email addresses, nor could we share email addresses received from Pôle emploi. Therefore, we used an anonymized matching algorithm to match addresses from both sources.

We present estimates of first stage equation [3] in Table 3. 0.2% of individuals in the control group created an

[27] In Appendix Tables A8 through A12, we show results obtained with robust standard errors and no clustering. The set of estimates which are statistically significant remains unchanged. These tables also report p-values from the Westfall-Young test to adjust for multiple testing: we test the null that the effects on all outcomes are equal to zero, for each table separately.

account on Bob Emploi (column 1) and, by design, none attended an informational session (column 2). By contrast, 13.2% of individuals in the treatment group created a Bob Emploi account and 20.5% attended an informational session. Taking both dimensions into account, the differential take-up between the treatment and control groups is 27.0 percentage points (column 3).

In the subsample of survey respondents, the differential take-up is higher, at 41.4 percentage points (column 4), reflecting the fact that individuals willing to attend an informational session and to try a new tool are also more willing to answer the survey. Column 5 checks the robustness of the first-stage results to using an alternative definition of take-up, which also sets B_i to 1 for survey respondents who report using Bob Emploi. (By construction, this alternative definition of take-up can only be applied to the survey subsample.) Indeed, it is possible that some actual users of Bob Emploi are missed by our procedure matching email addresses, for instance because the same person used different addresses for Bob Emploi and Pôle emploi. On the other hand, some respondents may mistakenly indicate that they use Bob Emploi, e.g., if they misunderstand which website the question refers to. The fraction of control individuals counted as treated under this alternative definition increases to 16.3% but, reassuringly, we obtain a differential take-up of 45.3 percentage points which is very close to the one reported in column 4. Therefore, our 2SLS regressions only define individuals as treated if they attended an informational session or if their email address was successfully matched with the address of a Bob Emploi user, including when we consider survey outcomes.

Finally, respondents were also surveyed about eight other websites providing job-search assistance, beyond Bob Emploi. The differential usage of Pôle emploi's platform "Emploi Store," the main other online resource which was mentioned during the informational sessions, is only 4.6 percentage points (see Appendix Table A4). Therefore, we can interpret our results as mostly driven by the usage of Bob Emploi.

4.3. Description of the takers

Table 4 compares the characteristics of takers and non-takers in the treatment group, along the same dimensions as those examined in Table 2. Note that given the very low take-up rate in the control group (0.2%), the vast majority of treatment group takers are compliers (not always-takers), i.e. individuals on whom the 2SLS effects are estimated.

Take-up was slightly larger among women, job seekers who had been unemployed for a longer time, and those with a university degree. The most striking pattern is that take-up markedly increases with age: job seekers older than 40 years old account for a much larger fraction of takers than non-takers (45.6% against 27.0%). This pattern is not due only to older people being more likely to attend informational sessions: it also holds if we focus on creations of Bob Emploi accounts (Appendix Table A5). Bob Emploi had more appeal for older job seekers, which is perhaps surprising as younger individuals tend to be more familiar with digital technologies.

4.4. Heterogeneity analysis

Benefits from using Bob Emploi may vary across groups of job seekers, and the fact that most of our average effects are null makes exploring impact heterogeneity particularly important. One cannot rule out that Bob Emploi helps certain groups while handicapping others, e.g., by wasting the time they spend searching for a job or by giving them bad advice. Education and age are the most natural dimensions of heterogeneity to consider. Younger and more educated people might find Bob Emploi easier to use and benefit more from it. Conversely, to the extent that Bob Emploi provides access to information that is also available on other platforms, it may be particularly beneficial for people who are less informed or initially less familiar with this type of tools.

Heterogeneity analysis by subgroups

To assess whether effects vary across age groups, we run the following specification:

$$Y_{it} = \alpha_1 + \beta_1 T_i + \sum_{k=1}^{3} \gamma_k D_i^k + \sum_{k=1}^{3} \delta_k T_i \times D_i^k + \mu_{i1}^{a,m} + \epsilon_{it1}, \tag{4}$$

where D_i^k's are dummies indicating whether the individual belongs to the subgroup k: less than 24 years old, 25 to 39 years old, or 40 to 54 years old. The δ_k coefficients measure differential effects in these groups relative to the omitted category of people older than 55.

We use the same specification to explore heterogeneous effects across job seekers with different levels of education. In that case, D_i^k's indicate no high-school diploma, vocational degree, and university diploma, and the omitted category is end-of-high-school diploma.

Generic machine learning

An important limitation of specifications in the form of equation [4] is that Bob Emploi may benefit individuals with profiles that are not captured by coarse age and education categories. Indeed, it is difficult to predict which groups of people benefit most from the customized advice provided by the website. At the very least, one would want to explore heterogeneity along other dimensions than just age and education. On the other hand, examining a large number of subgroups implies repeatedly testing the significance of δ_k coefficients and, thus, possibly making false discoveries. We use the generic machine learning approach developed by Chernozhukov et al. (2018) to explore heterogeneous effects in a more agnostic way, i.e. without imposing any assumption about the relevant dimensions of heterogeneity, while also avoiding the pitfalls of multiple hypothesis testing.

We start by randomly splitting our dataset in half between a training dataset and an estimating dataset. We use the training dataset to predict the effect of the treatment based on the following individual covariates as well as their interactions: age, gender, number of children, years spent working in the sector of interest, whether the job seeker looks for a part-time job or not, geographic indicators (dummies for each département and for whether the individual lives in a sensitive urban area), country of citizenship, reason why the job seeker is unemployed (end of a short term contract, end of an internship, contractual termination of a long term contract, end of a temporary contract, resignation, dismissal on economic grounds, other dismissals, others motives), availability to take a new job (indicating whether the job seeker is immediately available for a new job, or whether they are not available because they are already employed, off sick, or for other reasons), target job sector, target type of contract (short or long term), level of qualification, type of support received from Pôle emploi (the level of assistance the job seeker receives from the caseworker), distance the job seeker is willing to travel between home and work, benefits received at baseline, and reservation wage.

We perform this prediction exercise using different algorithms: gradient boosting, random forest, elastic net, causal forest with and without local centering, and Rlearner techniques combined with boosting and with lasso (see Appendix A.3 for more details). We define the resulting prediction function as $\widehat{\text{CATE}}(X)$ (for Conditional Average Treatment Effect). We also estimate the baseline outcome in the control group, using the same set of covariates: $\hat{B}(X) = \mathbb{E}\big[Y(0)|X\big]$.

Following Chernozhukov et al. (2018), we then use the estimating dataset to compute unbiased estimates. Our first regression aims to detect whether heterogeneous treatment effects are present or not:

$$Y_{it} = \theta_1 \hat{B}(X_i) + \theta_2 \widehat{\text{CATE}}(X_i) + \beta^{\text{ITT}}\big(T_i - p(X_i)\big) + \beta^{\text{HET}}\big(T_i - p(X_i)\big)\Big(\widehat{\text{CATE}}(X_i) - \overline{\text{CATE}}_M\Big) + \mu_{i1}^{a,m} + \epsilon_{it1}, \quad (5)$$

where $p(X_i)$ is the propensity score, $\overline{\text{CATE}}_M = \frac{1}{\#\{i \in M\}} \sum_{i \in M} \widehat{\text{CATE}}(X_i)$ is the mean predicted impact on sample M used for the regressions, and we use weights $w(X_i) = \frac{1}{p(X_i)(1-p(X_i))}$. β^{ITT} measures the mean treatment effect. The coefficient of interest is β^{HET}: it is different from zero if $\widehat{\text{CATE}}(X_i)$ captured heterogeneous effects.

Our second regression estimates Sorted Group Average Treatment Effects (GATES), which correspond to quintiles of estimated impact, and it enables us to compare them.

5. Results

We estimate the effects of Bob Emploi on four families of outcomes: job-search effort, scope, and expectations; job-search methods; well-being; and employment. In all tables, we show ITT effects in Panel A and 2SLS effects in Panel B. All effects reported in the text are ITT.

5.1. Job-search effort, scope, and expectations

We first investigate how Bob Emploi affects job seekers' reemployment expectations, measured using survey data, in Table 5.

Expectations of Bob Emploi users may improve if they find the advice provided by the website useful. On the other hand, information available on the website may make them realize that their reemployment chances are lower than expected. In the control group, 57.4% of respondents believe that they will find a job within the next three months. The difference between the treatment and control groups is small and nonsignificant. We can thus rule out that Bob Emploi demotivated job seekers. Unfortunately, we also rule out that Bob Emploi made job seekers more realistic about their chance of finding a job, despite the diagnosis of employability provided to them. Comparing respondents' expectations with their actual outcomes, measured using administrative data, we find that 69.3% of job seekers in the control group were overly optimistic – in the sense that it took them more time than they expected to find a job –, 12.3% overly pessimistic, and 18.4% realistic. None of these fractions is significantly different between the treatment and control groups (columns 7 to 9).

We now turn to effects on job seekers' search effort, and the types of jobs they apply for, reported in Table 6. Outcomes are measured using survey data (in columns 1 to 3) and administrative data (in columns 4 to 7).

Similarly as effects on expectations, effects on effort are ambiguous ex ante. By giving users concrete tips to improve their job search and helping them set reachable targets, Bob Emploi may increase the returns to effort. In turn, this may either increase users' motivation and the time they spend searching for a job or decrease that time, if users have quantitative targets in mind (e.g., in terms of number of interviews received) and they realize that they now need less time to achieve them. On net, we do not observe any effect on the number of hours respondents declare spending on their job search: this time is just over seven hours and a half in both the control and treatment groups (column 1). Survey respondents in the treatment group are not more likely to report sending at least one unsolicited application per week either (column 2). Using administrative data, we detect a consistent nonsignificant effect on the total number of online applications sent through Pôle emploi's job board and initiated by job seekers in the 18 months after the intervention (column 4). We also obtain small and nonsignificant effects on the total number of applications and on the subset of applications initiated by caseworkers and by employers, whether one measures these outcomes six months or 18 months after the intervention (Appendix Tables A6 and A7).

When a job seeker is targeting an occupation or an area where the labor market is particularly tight, Bob Emploi's algorithm often encourages them to consider expanding their search to related occupations or neighboring areas. However, we do not observe any significant effect on search scope. Compared to the control group, survey respondents in the treatment group are not more likely to report applying to jobs located outside the municipality where they live (column 3), and effects on the number of applications sent by job seekers through Pôle emploi for jobs outside their municipality and outside their target sector are both small and nonsignificant, with negative signs (columns 5 and 6). The treatment also fails to increase job seekers' propensity to apply for jobs advertising a salary lower than the reservation wage they indicated at the beginning of their unemployment spell (column 7).

5.2. Job-search methods

In addition to offering job seekers a diagnosis of reemployment perspectives, Bob Emploi provides them with actionable advice on how to improve their search strategy. Some of the most frequent recommendations include leveraging their personal and professional networks, adapting their application to the specifics of each job listing, and following up with recruiters after sending off applications. Furthermore, Bob Emploi often refers job seekers to other resources provided by Pôle emploi or other private websites. Columns 1 to 7 of Table 7 use our survey data to investigate whether Bob Emploi changed job-search methods along these different dimensions. A majority of respondents in the control group report using their personal and professional networks when looking for a job. Effects on both outcomes are positive and significant at the 1 and 10% levels, respectively (columns 1 and 2). The 1.7 and 0.8 percentage point increases correspond to 2.7% and 1.1% of the means in the control group. Job seekers in the treatment group are not more likely to report adapting their CV and cover letters to job listings (columns 3 and 4), but we do find a small positive effect on following up with recruiting firms. The point estimate is equal to 1.0 percentage point (2.4% of the control mean) and it is significant at the 5% level (column 5). We also observe a 4.7% (0.065/1.369, column 6) increase in the number of Pôle emploi websites used by job seekers, which is significant at the 1% level. This effect may be driven both by using Bob Emploi and by the short presentation of Pôle emploi's resource "Emploi Store" during the informational sessions. By contrast, the effect on the number of private websites used by job seekers is small, negative, and nonsignificant (column 7).

Columns 8 and 9 turn to the use of public job-search assistance offline services, as recorded in Pôle emploi's administrative data. A possible concern is if private websites such as Bob Emploi discourage job seekers from engaging in conversations with their caseworkers or from participating in the programs offered by Pôle emploi, replacing those with a more accessible but less effective form of assistance. An alternative view is that Bob Emploi and Pôle emploi are complements, not substitutes. Given the close partnership between Bayes Impact and Pôle emploi, reflected in the content of the advice provided by Bob Emploi, one may even expect the treatment to increase the use of Pôle emploi's offline services. If anything, our results support the latter prediction of complementarity between services. We do not observe any significant impact on the number of workshops attended at Pôle emploi (column 9), and the treatment *increased* the fraction of job seekers who had at least one meeting with their caseworker within the six months following the intervention by 2.4 percentage points (4.2%, column 8). An important caveat is that this effect may be partly driven by participation in the informational sessions, which created an opportunity for job seekers to meet with their caseworker.

5.3. Well-being and life balance

Table 8 reports effects on the self-reported well-being of survey respondents. Along with a direct loss of income and the arduous prospect of having to search for a job, unemployment is often associated with loss of a sense of meaning and dislocation of social bounds (McKee-Ryan et al. 2005). Unemployed workers' unhappiness may further increase if their job search is unsuccessful (Krueger and Mueller 2011), which may in turn lower their motivation and effort, creating a vicious circle. Individuals who have stayed unemployed for too long may eventually get discouraged and drop out of the labor market. By providing job seekers with actionable advice to optimize their search efforts in a user-friendly manner, Bob Emploi may increase their sense of having a purpose and being supported. Of course, well-being could further increase if using Bob Emploi improved actual search outcomes.

In the control group, individuals rate their well-being slightly above 5, on a scale from 0 to 10 (column 1), and they are more likely to indicate that they feel motivated than supported in their job search (columns 2 and 3). This discrepancy emphasizes the importance of improving the assistance available to job seekers. With treatment, we do observe a small increase in feeling supported by 1.4% (0.052/3.833), which is significant at the 10% level, but also a *decrease* in feeling motivated by 0.7% (0.047/6.958), which is significant at the 5% level. The effect of the treatment on overall well-being is close to zero and nonsignificant. The survey further included three questions recording whether job seekers participated in sport, arts, and cultural or community activities at least once a month. Column 4 uses as outcome a dummy equal to 1 if the respondent participates in any of these. Although Bob Emploi often advises its users to keep a balanced life and get involved in social activities, we do not observe any significant difference between the treatment and the control groups.

5.4. Employment outcomes

As interesting and important the outcomes explored to this point are, Bob Emploi's ultimate objective is to help unemployed workers find a job. In Table 9, we report effects on four individual-level outcomes, all measured using administrative data: experiencing any employment episode following the intervention, obtaining a stable job (a contract of more than six months), total number of days in unemployment, and total amount of unemployment benefits.[28] Columns 1 to 4 and 5 to 8 report effects for six months and 18 months following the intervention, respectively. Our outcomes accommodate the fact that some individuals find a job and then go back to unemployment, sometimes with multiple switches in and out of employment. For such individuals, the outcome "experiencing any employment episode" is equal to 1, and we take all unemployment periods into account when computing the number of days in unemployment and unemployment benefits.

We find no significant effects on getting a job and on the total number of days in unemployment at six months or 18 months after the intervention. These null results are precisely estimated. For instance, considering the upper bound of the 95 percent confidence interval, we can reject any effect higher than 0.4 percentage points on experiencing some employment episode within the 18 months following the intervention, even though 40.9% of the sample has still not found any job by that time (column 5). Similarly, we can reject any reduction of the average time spent unemployed by more than two days within the 18 months post intervention (column 7). The lack of effect on finding a job does not preclude the possibility of Bob Emploi users finding *better* jobs. However, the effect on obtaining a stable job is also close to zero and nonsignificant. Given the lack of effect on employment, the effect on the total benefits received during unemployment is, unsurprisingly, also null.

Focusing on the six- and 18-month mark following the intervention is somewhat arbitrary. Therefore, Figures

[28] Pôle emploi's data record all employment episodes, except for the following cases: public sector jobs using some types of contracts, jobs abroad, interim work of unknown length, individuals hired by a private individual, and individuals starting their own company.

2 and 3 complete the results shown in Table 9 by plotting the four same outcomes over time, separately in the treatment and control groups (top graphs), as well as the treatment effects and their confidence interval (bottom graphs). All level series are concave, in line with models of job-search frictions and other empirical evidence. These figures bring additional support for the conclusion that Bob Emploi did not improve employment outcomes. Indeed, the series representing outcomes in the treatment and control groups are nearly indistinguishable from each other, and all effects oscillate around zero, with tight confidence intervals.

5.5. Heterogeneity analysis

Lastly, we ask whether the null average effects of Bob Emploi on employment outcomes hide positive effects for some groups of job seekers, while other groups are not affected or affected negatively. Our exploration of heterogeneity focuses on the total duration of unemployment in the six or 18 months following the intervention.[29]

Effects by age and education

We first estimate specifications in the form of equation [4], in which we interact the treatment dummy with age groups or education levels. The point estimate on the treatment dummy corresponds to the effect on the omitted group (people older than 55 and people with an end-of-high-school diploma, respectively) and the coefficients on the interactions measure the differential effect on the corresponding groups. We also report the point estimates and standard errors of linear combinations of the treatment with each of the interactions to test whether effects on any group are different from zero. As shown in Table 10, all coefficients and all their combinations are nonsignificant, whether we measure effects on total unemployment duration six or 18 months after the intervention. We infer that effects are null for all age groups and education levels and that they cannot be distinguished from each other.

Generic machine learning results

Results of the generic machine learning approach are reported in Table 11 and Appendix Figure A5 (for effects after six months) and in Table 12 and Appendix Figure A6 (for effects after 18 months).

We first report results of specifications in the form of equation [5], using the two best algorithms chosen based on Chernozhukov et al. (2018)'s criteria. The best approximations of individualized impacts are given by random forest and causal forest with local centering, for effects after six months, and by boosting and causal forest without local centering, for effects after 18 months. As shown in Table 11, Panel A, the estimate of the mean effect, β^{ITT}, is nonsignificant, consistent with the lack of average effect on unemployment duration already visible in Table 9, column 3. Moreover, we cannot reject that $\beta^{\text{HET}} = 0$, meaning that we do not detect heterogeneity. Results shown in Table 12, Panel A, are qualitatively similar.

[29] We explore heterogeneous effects on this continuous outcome rather than the dummy equal to 1 for experiencing some employment episode because the generic machine learning method includes a linear regression phase which does not work as well with binary variables.

In principle, our inability to detect heterogeneity may be due to bad approximations of the individual treatment effect by the algorithms, low statistical power due to sample splitting, or there being truly no heterogeneity. The first of these explanations is unlikely to hold since we obtain consistent results when using different algorithms, each relying on different assumptions. Loss of statistical power due to sample splitting is also unlikely to be a problem here given our very large sample size and our high take-up. Therefore, we conclude that our results reflect the true absence of heterogeneity.

Second, Appendix Figures A5 and A6 report GATES estimates, corresponding to quintiles of estimated impact, and Panel B's of Tables 11 and 12 compare the first and fifth quintiles. No matter what machine learning algorithm we use, we find that none of the quintiles of treatment effects are statistically significant and that the difference between the top and bottom quintiles is not significant either. These results bring additional support for our conclusion that the intervention did not affect unemployment duration for any group of individuals.

6. Conclusion

Persistently high rates of unemployment in France and in other OECD countries and the ensuing economic, social, and psychological costs have created a demand by policymakers, job seekers, and the broader public, for new types of interventions. In this paper, we report the results of a large-scale randomized experiment evaluating the impact of Bob Emploi, a private website using data analysis to provide job-search assistance to the unemployed. We implemented the experiment in partnership with the organization that created the website and the French public employment services, giving us access to high-quality individual-level administrative data and resulting in a very large sample size.

We conclude that the website does not significantly reduce unemployment. Job seekers in the treatment group do not adjust their reemployment expectations and they remain as overly optimistic as those in the control group, despite the individual diagnosis of employability provided by Bob Emploi. Treated individuals become slightly more likely to use their personal and professional network when searching for a job and to follow up with recruiting firms after sending their application. However, the effects on sending unsolicited applications and on the total number of applications sent through Pôle emploi's job board are small and nonsignificant, and individuals in the treatment group do not widen the scope of their search, despite the frequent invitations of the website to explore jobs in locations or sectors where the labor market is less tight. Overall, the modest effects we find on search practices increase neither unemployed workers' likelihood to find a job nor their overall self-reported well-being. Our null effects on the likelihood to find a job and on the duration of unemployment after six or 18 months are precisely estimated, given our large sample and our strong first stage. Furthermore, we do not find any effect on unemployment for any group, whether we specifically consider groups defined by age or education or use the more agnostic and comprehensive generic machine learning approach developed by Chernozhukov et al. (2018).

Beyond Bob Emploi, a growing number of private websites offer job-search assistance services to unemployed workers. If anything, Bob Emploi may have been in a better position than other websites to have a positive impact. Indeed, it was built by experienced developers, received significant public funding, and obtained access to rich labor market data from the French public employment services. We cannot exclude that better data and algorithms, coupled with dynamic matching functionalities of job boards, may be more effective. Yet, our results make it clear that such improvements are not low-hanging fruits. Overall, this study suggests that the enthusiasm around the potential for digital platforms like Bob Emploi to help reduce unemployment should be toned down.

Table 1: Comparison between Bob Emploi and other websites providing job search assistance

	Country	Job listings	Personalized advice	General advice	Planning assistance	Motivational content
Bob Emploi	France		x	x	x	x
Meteojob	France	x		x		
Jobijoba	France	x	x	x		
Monkey tie	France	x	x			
MEMO	France				x	
Indeed	U.S.	x		x		x
ZipRecruiter	U.S.	x	x			
Snagajob	U.S.	x	x	x	x	

Notes: This table compares the functionalities of Bob Emploi with the features of other websites providing job search assistance, in France and in the U.S.

Table 2: Summary statistics

	Control group mean (1)	Regression coefficient on treatment dummy (2)
Female	50.719	-0.131
	[49.995]	(0.236)
Age		
Under 25 year old	23.110	-0.290
	[42.154]	(0.173)*
Between 25 and 39 year old	45.418	0.262
	[49.790]	(0.215)
Between 40 and 54 year old	23.331	0.071
	[42.294]	(0.192)
Above 55 year old	8.141	-0.043
	[27.347]	(0.127)
Seniority in unemployment		
Less than 3 months	25.246	0.099
	[43.442]	(0.181)
Between 3 and 6 months	21.306	-0.204
	[40.947]	(0.179)
More than 6 months	53.203	0.104
	[49.898]	(0.222)
Level of education		
No high-school diploma	16.185	-0.030
	[36.832]	(0.169)
Vocational degree	30.598	-0.113
	[46.082]	(0.189)
End-of-high-school diploma	18.664	0.140
	[38.962]	(0.162)
University degree	34.552	0.004
	[47.554]	(0.197)
N	92752	212277

Notes: This table reports summary statistics and balance checks for a set of pre-intervention covariates. Column 1 shows mean values in the control group, with standard deviations in brackets. Column 2 reports regression coefficients on the treatment dummy, with standard errors in parentheses. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively.

Table 3: First stage

	Bob account	Attended meeting	Bob account + attended meeting	Bob account + attended meeting	Bob account + attended meeting + reported use
	(1)	(2)	(3)	(4)	(5)
Treatment	0.130	0.205	0.270	0.414	0.453
	(0.002)***	(0.006)***	(0.006)***	(0.007)***	(0.006)***
Control mean	0.002	0.000	0.002	0.004	0.163
N	212277	212277	212277	47327	47327
Sample	All	All	All	Survey	Survey

Notes: This table shows first stage results from equation [3]. The outcome is a dummy equal to 1 for individuals who created an account on Bob Emploi (in column 1), who attended an informational session (in column 2), who created an account or attended an informational session (in columns 3 and 4), and who created an account, attended an informational session, or reported using Bob Emploi in the survey (in column 5). The sample is restricted to survey respondents in columns 4 and 5. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively.

Table 4: Differences between takers and non-takers in the treatment group

	Non-taker mean (1)	Regression coefficient on taker dummy (2)
Female	49.496	4.395
	[49.998]	(0.342)***
Age		
Under 25 year old	26.160	-12.295
	[43.951]	(0.316)***
Between 25 and 39 year old	46.845	-6.340
	[49.901]	(0.383)***
Between 40 and 54 year old	20.706	10.852
	[40.520]	(0.313)***
Above 55 year old	6.290	7.783
	[24.278]	(0.243)***
Seniority in unemployment		
Less than 3 months	25.517	-1.437
	[43.596]	(0.310)***
Between 3 and 6 months	21.746	-1.644
	[41.252]	(0.291)***
More than 6 months	52.509	3.066
	[49.937]	(0.352)***
Level of education		
No high-school diploma	16.470	-1.631
	[37.091]	(0.240)***
Vocational degree	31.500	-1.587
	[46.452]	(0.327)***
End-of-high-school diploma	19.587	-2.512
	[39.687]	(0.273)***
University degree	32.443	5.731
	[46.816]	(0.373)***
N	87015	119525

Notes: This table reports summary statistics for individuals in the treatment group who did not take up the treatment and the results of tests of equal means with the takers for a set of pre-intervention covariates. Column 1 shows mean values for non-takers, with standard deviations in brackets. Column 2 reports regression coefficients on the take-up dummy, with standard errors in parentheses. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively.

Table 5: Impact on reemployment expectations

	Expects to find a job in less than a month	Expects to find a job in 1 to 3 months	Expects to find a job in 4 to 6 months	Expects to find a job in 7 to 12 months	Expects to find a job in more than a year	Expects to not find a job	Overly optimistic	Overly pessimistic	Realistic
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: ITT estimation									
Treatment	0.001	-0.007	0.006	-0.001	-0.003	0.003	0.006	0.000	-0.006
	(0.005)	(0.006)	(0.005)	(0.004)	(0.003)	(0.004)	(0.006)	(0.004)	(0.005)
Panel B: Instrumental variable estimation: *taker* instrumented with *treatment*									
Taker	0.002	-0.016	0.014	-0.001	-0.006	0.007	0.013	0.000	-0.013
	(0.012)	(0.014)	(0.012)	(0.009)	(0.006)	(0.008)	(0.014)	(0.009)	(0.012)
Control mean	0.185	0.389	0.200	0.098	0.050	0.078	0.693	0.123	0.184
N	23829	23829	23829	23829	23829	23829	23829	23829	23829

Notes: Panel A shows ITT results from equation [1] and Panel B shows 2SLS results from equation [2]. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. All regressions use dependent variables from survey data. Overly optimistic job seekers are job seekers that found a job later than what they anticipated. Overly pessimistic job seekers are job seekers that found a job earlier than anticipated. Realistic job seekers are job seekers that found within their anticipated timeline.

Table 6: Impact on effort and scope of job search

	Survey data			Administrative data (18 months)			
	Hours per week spent on job search	At least one unsolicited application per week	Applies to jobs outside their municipality	Number of online applications	Number of online applications outside their municipality	Number of online applications outside their target sector	Number of online applications below their reservation wage
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: ITT estimation							
Treatment	-0.083	0.004	0.000	-0.019	-0.016	-0.040	-0.005
	(0.068)	(0.005)	(0.005)	(0.099)	(0.092)	(0.061)	(0.049)
Panel B: Instrumental variable estimation: *taker* instrumented with *treatment*							
Taker	-0.197	0.010	0.000	-0.071	-0.060	-0.150	-0.019
	(0.161)	(0.011)	(0.012)	(0.368)	(0.342)	(0.225)	(0.183)
Control mean	7.624	0.458	0.505	2.930	2.331	2.112	1.170
N	42985	41855	39117	212277	212277	212277	212277

Notes: Panel A shows ITT results from equation [1] and Panel B shows 2SLS results from equation [2]. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. Regressions in columns 1 to 3 use dependent variables from survey data. Regressions in columns 4 to 7 use dependent variables from administrative data from Pôle emploi, which are available for the full sample. Outcomes in these columns are defined for the period of 18 months after the intervention and only count applications initiated by job seekers.

Table 7: Impact on job search methods

	Survey data							Administrative data	
	Uses personal network	Uses professional network	Adapts CV to job listing	Adapts cover letter to job listing	Follows up with recruting firms	Number of Pôle emploi websites used	Number of private websites used	At least one meeting with caseworker	Number of workshops attended at Pôle emploi
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: ITT estimation									
Treatment	0.017	0.008	0.006	-0.003	0.010	0.065	-0.006	0.024	-0.002
	(0.005)***	(0.005)*	(0.005)	(0.004)	(0.005)**	(0.008)***	(0.011)	(0.003)***	(0.010)
Panel B: Instrumental variable estimation: *taker* instrumented with *treatment*									
Taker	0.040	0.019	0.014	-0.007	0.023	0.154	-0.013	0.090	-0.006
	(0.012)***	(0.011)*	(0.012)	(0.010)	(0.012)**	(0.020)***	(0.026)	(0.012)***	(0.036)
Control mean	0.623	0.755	0.608	0.745	0.421	1.369	1.550	0.572	0.653
N	39776	37566	39538	40011	39618	41448	41289	212277	212277

Notes: Panel A shows ITT results from equation [1] and Panel B shows 2SLS results from equation [2]. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. Regressions in columns 1 to 7 use dependent variables from survey data. Regressions in columns 8 and 9 use dependent variables from administrative data from Pôle emploi, which are available for the full sample. Outcomes in these columns are defined for the period of 18 months after the intervention.

Table 8: Impact on well-being

	Overall well-being	Feels motivated during job search	Feels supported during job search	Partici-pates in non-job-related activities
	(1)	(2)	(3)	(4)
Panel A: ITT estimation				
Treatment	0.026	-0.047	0.052	0.003
	(0.022)	(0.022)**	(0.029)*	(0.004)
Panel B: Instrumental variable estimation: *taker* instrumented with *treatment*				
Taker	0.062	-0.112	0.126	0.008
	(0.052)	(0.053)**	(0.070)*	(0.011)
Control mean	5.167	6.958	3.833	0.772
N	44724	44724	44724	37035

Notes: Panel A shows ITT results from equation [1] and Panel B shows 2SLS results from equation [2]. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. All regressions use dependent variables from survey data.

Table 9: Impact on employment outcomes

	6 months after the intervention				**18 months after the intervention**			
	Employment episode	Stable employment episode	Unemployment duration	Unemployment benefits	Employment episode	Stable employment episode	Unemployment duration	Unemployment benefits
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: ITT estimation								
Treatment	0.011	-0.025	-0.467	2.642	0.022	0.094	-0.382	12.667
	(0.230)	(0.149)	(0.293)	(19.577)	(0.213)	(0.187)	(0.833)	(40.728)
Panel B: Instrumental variable estimation: *taker* instrumented with *treatment*								
Taker	0.039	-0.093	-1.732	9.791	0.081	0.350	-1.415	46.940
	(0.853)	(0.554)	(1.092)	(72.534)	(0.791)	(0.692)	(3.090)	(150.895)
Control mean	40.485	15.163	199.260	3031.994	59.128	25.578	418.875	6069.893
N	212277	212277	212277	212277	212277	212277	212277	212277

Notes: Panel A shows ITT results from equation [1] and Panel B shows 2SLS results from equation [2]. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. All regressions use dependent variables from administrative data. *Employment episode* is a dummy equal to 1 for job seekers who experienced at least one employment episode following the intervention. *Stable employment episode* is a dummy equal to 1 for job seekers who experienced at least one stable employment episode (long-term contract, or short-term contract for more than 6 months) following the intervention. *Unemployment duration* counts the total number of days in unemployment since the intervention. *Unemployment benefits* measures the total benefits that job seekers received following the intervention. Regressions in columns 1 to 4 use outcomes measured six months after the intervention. Regressions in columns 5 to 8 use outcomes measured 18 months after the intervention.

Table 10: Impact on unemployment duration, heterogeneity by age and education

	Unemployment duration 6 months after the intervention (1)	Unemployment duration 18 months after the intervention (2)	Unemployment duration 6 months after the intervention (3)	Unemployment duration 18 months after the intervention (4)
Treatment [T]	-0.493	-1.323	-0.417	-1.364
	(0.445)	(1.187)	(0.707)	(1.900)
Treatment x Under 25 yo [T1]	0.224	2.548		
	(0.826)	(2.101)		
Treatment x 25-39 yo [T2]	-0.0513	0.402		
	(0.727)	(2.124)		
Treatment x 40-54 yo [T3]	-0.643	1.998		
	(1.051)	(3.314)		
Treatment x No high-school diploma [T1]			0.131	0.927
			(1.019)	(2.758)
Treatment x Vocational degree [T2]			-0.107	1.925
			(0.879)	(2.423)
Treatment x University degree [T3]			-0.0988	0.757
			(0.960)	(2.596)
Control mean	199.260	418.875	199.260	418.875
N	212277	212277	212277	212277
Linear combination: [T] + [T1] = 0	-.268	1.225	-.286	-.437
	(.676)	(1.737)	(.728)	(2.065)
Linear combination: [T] + [T2] = 0	-.544	-.922	-.524	.561
	(.572)	(1.78)	(.543)	(1.556)
Linear combination: [T] + [T3] = 0	-1.135	-.922	-.516	-.607
	(.966)	(1.78)	(.535)	(1.519)

Notes: This table shows ITT results from equation [4]. Regressions in columns 1 and 3 (resp. 2 and 4) use as outcomes the total number of days in unemployment in the six months (resp. 18 months) after the intervention. Regressions in columns 1 and 2 (resp. 3 and 4) measure treatment impact heterogeneity by age (resp. education level). All regressions include agency x month fixed effects. Standard errors in parentheses are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level respectively.

Table 11: Impact on unemployment duration at 6 months, generic machine learning

(a) β^{ITT} and β^{HET} estimates from equation [5]

	Random forest		**Causal forest with local centering**	
	β^{ITT}	β^{HET}	β^{ITT}	β^{HET}
Estimate	-0.57	0.02	-0.57	-0.08
Confidence interval (90%)	[-1.35, 0.21]	[-0.03, 0.08]	[-1.36, 0.22]	[-0.39, 0.25]
Adjusted p-value	(0.30)	(0.82)	(0.32)	(1.00)

(b) GATES and difference between top and bottom quintile treatment effects

	Random forest			**Causal forest with local centering**		
	Q5	Q1	Difference	Q5	Q1	Difference
Estimate	0.03	-1.13	1.16	-0.99	-0.44	-0.68
Confidence interval (90%)	[-1.92, 1.86]	[-3.07, 0.78]	[-1.59, 3.91]	[-2.82, 0.85]	[-2.31, 1.47]	[-3.35, 1.93]
Adjusted p-value	(1.00)	(0.50)	(0.39)	(0.60)	(1.00)	(0.52)

Notes: Quintiles Q5 and Q1 are estimations of γ_5 and γ_1 from equation [7] (in Appendix A.3). We also test whether the difference between both is significantly different from 0. Following Chernozhukov et al. (2018), all displayed estimations are medians over 100 splits in half, confidence intervals are adjusted, and so are p-values. Standard errors are clustered at the agency level.

Table 12: Impact on unemployment duration at 18 months, generic machine learning

(a) β^{ITT} and β^{HET} estimates from equation [5]

	Boosting		**Causal forest without local centering**	
	β^{ITT}	β^{HET}	β^{ITT}	β^{HET}
Estimate	-0.23	-0.06	-0.46	-0.11
Confidence interval (90%)	[-2.51, 2.03]	[-0.19, 0.07]	[-2.70, 1.83]	[-0.37, 0.17]
Adjusted p-value	(1.00)	(0.75)	(1.00)	(0.81)

(b) GATES and difference between top and bottom quintile treatment effects

	Boosting			**Causal forest without local centering**		
	Q5	Q1	Difference	Q5	Q1	Difference
Estimate	-1.80	0.75	-2.10	-1.63	0.97	-2.07
Confidence interval (90%)	[-6.81, 3.12]	[-4.41, 5.87]	[-9.18, 4.92]	[-6.78, 3.58]	[-4.43, 6.14]	[-9.43, 5.21]
Adjusted p-value	(0.84)	(1.00)	(0.40)	(0.90)	(1.00)	(0.46)

Notes: Quintiles Q5 and Q1 are estimations of γ_5 and γ_1 from equation [7] (in Appendix A.3). We also test whether the difference between both is significantly different from 0. Following Chernozhukov et al. (2018), all displayed estimations are medians over 100 splits in half, confidence intervals are adjusted, and so are p-values. Standard errors are clustered at the agency level.

Figure 1: Timeline of the experiment

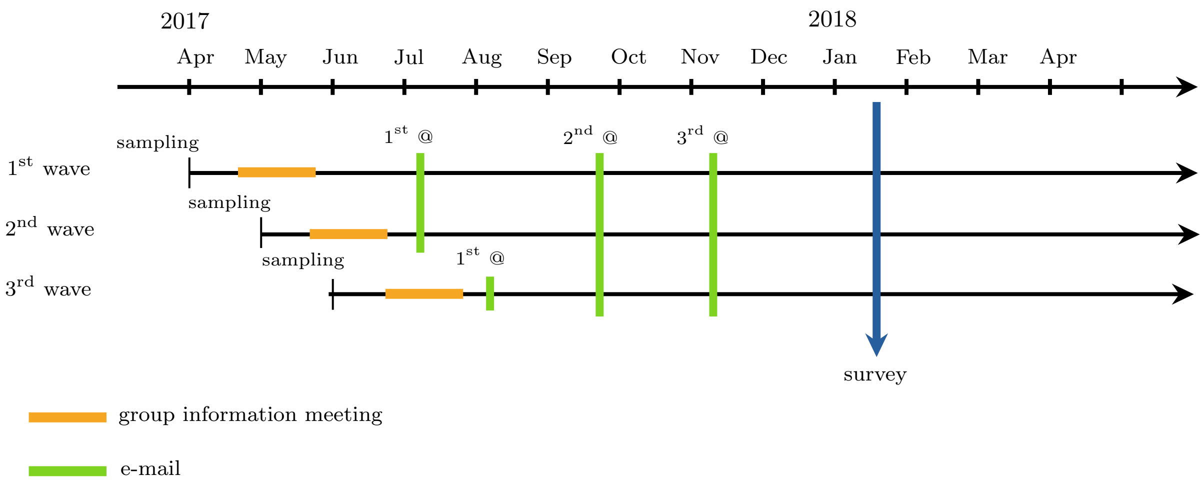

Notes: The figure shows the timeline of the experiment. The sample was drawn in three separate installments on April 1, May 1, and June 1. Informational sessions took place between late April and late May (resp. late May and late June, and late June and late July) for individuals drawn on April 1 (resp. May 1 and June 1). Three reminder emails were sent out on July 7 and 26 (for individuals drawn in the first two months and in the third month, respectively), September 28, and November 13. The online survey was administered on January 23, 2018.

Figure 2: Impact on employment episodes

(a) Employment episode

(b) Stable employment episode

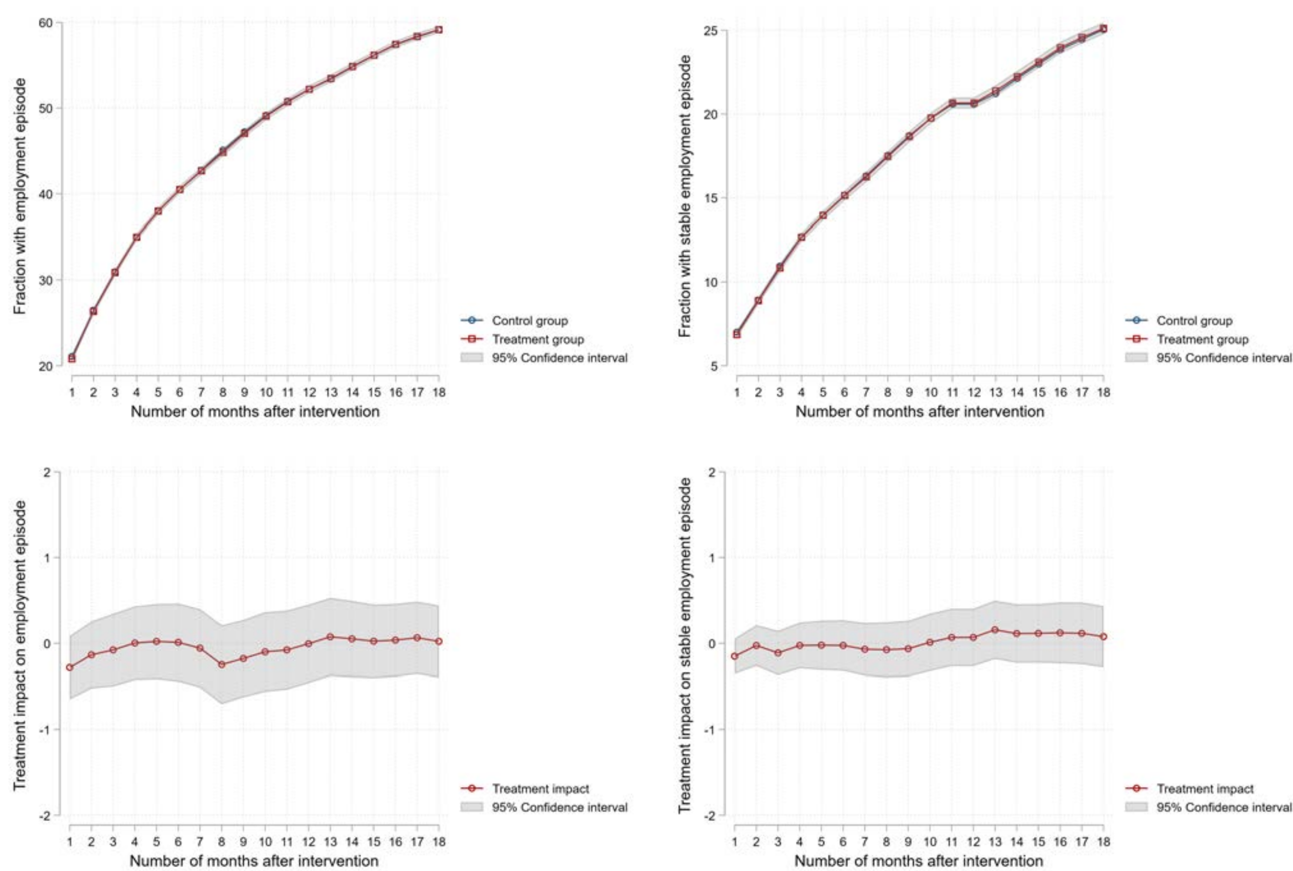

Notes: Figure 2a (resp. 2b) plots the fraction of individuals who have experienced any employment episode (resp. any stable employment episode) one month to 18 months after the intervention, separately in the control and treatment groups (top graph), as well as ITT results from equation [1] (bottom graph).

Figure 3: Impact on unemployment duration and unemployment benefits

(a) Unemployment duration

(b) Unemployment benefits

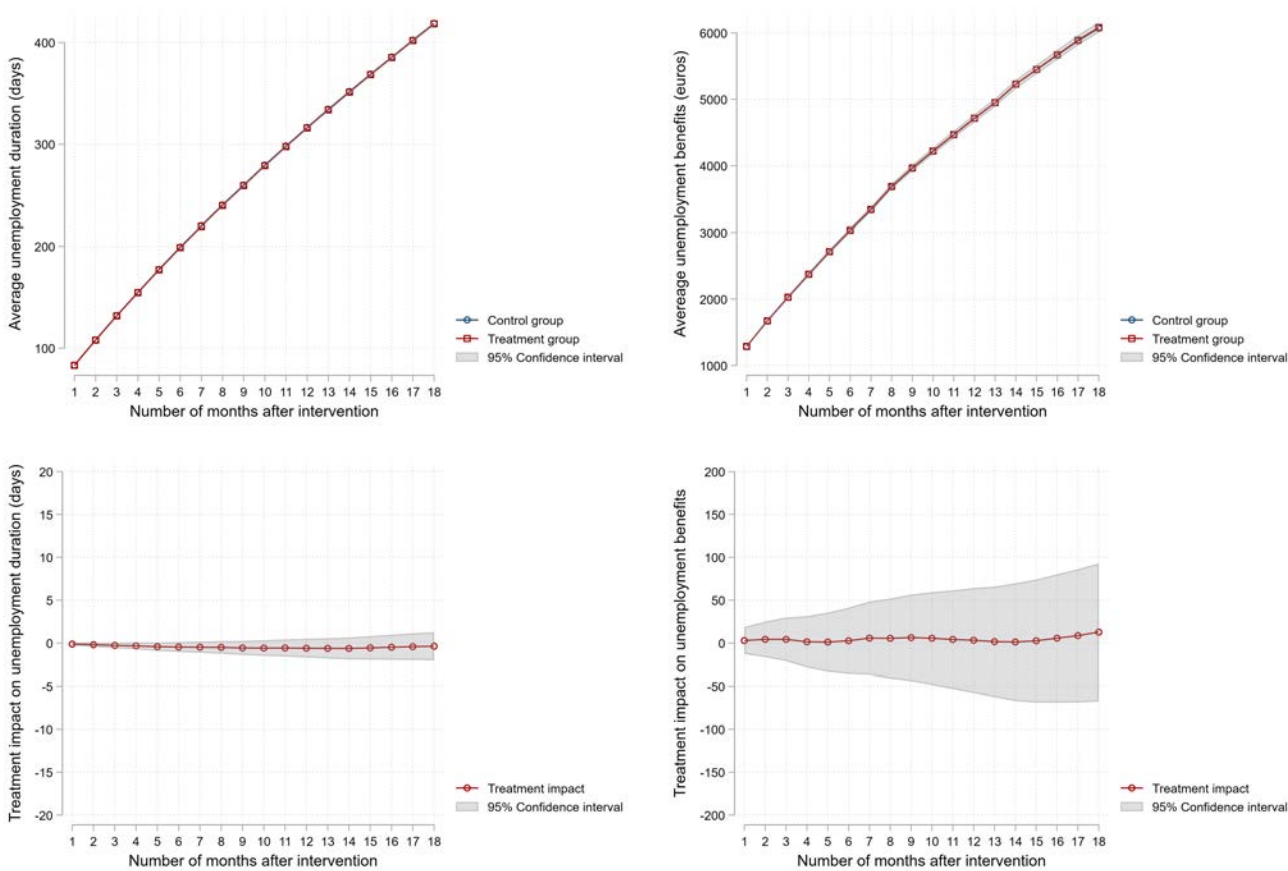

Notes: Figure 3a (resp. 3b) plots the mean total number of days in unemployment (resp. mean total amount of unemployment benefits) one month to 18 months after the intervention, separately in the control and treatment groups (top graph), as well as ITT results from equation [1] (bottom graph).

References

Abel, Martin, Rulof Burger, Eliana Carranza, and Patrizio Piraino. 2019. "Bridging the Intention-Behavior Gap? The Effect of Plan-Making Prompts on Job Search and Employment". *American Economic Journal: Applied Economics* 11 (2): 284–301.

Altmann, Steffen, Armin Falk, Simon Jäger, and Florian Zimmermann. 2018. "Learning About Job Search: A Field Experiment with Job Seekers in Germany". *Journal of Public Economics* 164:33–49.

Athey, Susan, Julie Tibshirani, Stefan Wager, et al. 2019. "Generalized Random Forests". *The Annals of Statistics* 47 (2): 1148–1178.

Autor, David. 2009. "Introduction to "Studies of Labor Market Intermediation"". In *Studies of Labor Market Intermediation*, 1–23. National Bureau of Economic Research, Inc.

Babcock, Linda, William J Congdon, Lawrence F Katz, and Sendhil Mullainathan. 2012. "Notes on Behavioral Economics and Labor Market Policy". *IZA Journal of Labor Policy* 1 (1): 1–14.

Banerjee, Abhijit V, Esther Duflo, and Rachel Glennerster. 2008. "Putting a Band-Aid on a Corpse: Incentives for Nurses in the Indian Public Health Care System". *Journal of the European Economic Association* 6 (2-3): 487–500.

Banerjee, Abhijit, Rukmini Banerji, James Berry, Esther Duflo, Harini Kannan, Shobhini Mukerji, Marc Shotland, and Michael Walton. 2017. "From Proof of Concept to Scalable Policies: Challenges and Solutions, with an Application". *Journal of Economic Perspectives* 31, no. 4 (): 73–102.

Belot, Michele, Philipp Kircher, and Paul Muller. 2019. "Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice". *The Review of Economic Studies* 86 (4): 1411–1447.

Breiman, Leo. 2001. "Random Forests". *Machine learning* 45 (1): 5–32.

Campbell, Carl M., and Kunal S. Kamlani. 1997. "The Reasons for Wage Rigidity: Evidence From a Survey of Firms". *The Quarterly Journal of Economics* 112 (3): 759–789.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Ivan Fernandez-Val. 2018. *Generic Machine Learning Inference on Heterogenous Treatment Effects in Randomized Experiments.* Tech. rep. National Bureau of Economic Research.

Cooper, Michael, and Peter Kuhn. 2020. "Behavioral Job Search". In *Handbook of Labor, Human Resources and Population Economics*, 1–22. Cham: Springer International Publishing.

DellaVigna, Stefano, Jörg Heining, Johannes Schmieder, and Dr. Simon Trenkle. 2020. *Evidence on Job Search Models from a Survey of Unemployed Workers in Germany.* IZA Discussion Papers 13189. Institute of Labor Economics (IZA).

DellaVigna, Stefano, and M Daniele Paserman. 2005. “Job Search and Impatience”. *Journal of Labor Economics* 23 (3): 527–588.

Gee, Laura K. 2019. “The More You Know: Information Effects on Job Application Rates in a Large Field Experiment”. *Management Science* 65 (5): 2077–2094.

Grossman, Guy, Macartan Humphreys, and Gabriella Sacramone-Lutz. 2020. “Information Technology and Political Engagement: Mixed Evidence from Uganda”. *The Journal of Politics* 82 (4): 1321–1336.

Herz, Benedikt, and Thijs van Rens. 2019. “Accounting for Mismatch Unemployment”. *Journal of the European Economic Association* 18 (4): 1619–1654.

Horton, John J. 2017. “The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment”. *Journal of Labor Economics* 35 (2): 345–385.

Krueger, Alan B, and Andreas I Mueller. 2012. “Time Use, Emotional Well-Being, and Unemployment: Evidence from Longitudinal Data”. *American Economic Review* 102 (3): 594–99.

Krueger, Alan, and Andreas Mueller. 2011. “Job Search, Emotional Well-Being and Job Finding in a Period of Mass Unemployment: Evidence from High-Frequency Longitudinal Data”. *Brookings Papers on Economic Activity* 42 (1 (Spring)): 1–81.

Kuhn, Peter, and Hani Mansour. 2014. “Is Internet Job Search Still Ineffective?” *The Economic Journal* 124 (581): 1213–1233.

Kuhn, Peter, and Mikal Skuterud. 2004. “Internet Job Search and Unemployment Durations”. *American Economic Review* 94 (1): 218–232.

Künzel, Sören R., Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. 2019. “Meta-learners for Estimating Heterogeneous Treatment Effects using Machine Learning”. *Proceedings of the National Academy of Sciences* 116 (10): 4156–4165.

Marinescu, Ioana, and Roland Rathelot. 2018. “Mismatch Unemployment and the Geography of Job Search”. *American Economic Journal: Macroeconomics* 10 (3): 42–70.

Marinescu, Ioana, and Daphné Skandalis. 2021. “Unemployment insurance and job search behavior”. *The Quarterly Journal of Economics* 136 (2): 887–931.

McCall, J. J. 1970. “Economics of Information and Job Search”. *The Quarterly Journal of Economics* 84 (1): 113–126.

McKee-Ryan, Frances, Zhaoli Song, Connie R Wanberg, and Angelo J Kinicki. 2005. “Psychological and Physical Well-Being During Unemployment: A Meta-Analytic Study”. *The Journal of Applied Psychology* 90 (1): 53–76.

Michaillat, Pascal. 2012. “Do Matching Frictions Explain Unemployment? Not in Bad Times”. *American Economic Review* 102 (4): 1721–1750.

Mortensen, Dale T. 1970. “A Theory of Wage and Employment Dynamics”. *Microeconomic Foundations of Employment and Inflation Theory*: 167–211.

Mueller, Andreas I., Johannes Spinnewijn, and Giorgio Topa. 2021. “Job Seekers’ Perceptions and Employment Prospects: Heterogeneity, Duration Dependence, and Bias”. *American Economic Review* 111, no. 1 (): 324–63.

Nie, X, and S Wager. 2020. “Quasi-Oracle Estimation of Heterogeneous Treatment Effects”. *Biometrika* 108 (2): 299–319.

Rubin, D.B. 1974. “Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies”. *Journal of Educational Psychology* 66 (5): 688–701.

Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2014. “Mismatch Unemployment”. *American Economic Review* 104 (11): 3529–64.

Skandalis, Daphné. 2019. “Breaking News: The Role of Information in Job Search and Matching”. *Working paper.*

Spinnewijn, Johannes. 2015. “Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs”. *Journal of the European Economic Association* 13 (1): 130–167.

Temin, Peter. 1990. “Socialism and Wages in the Recovery from the Great Depression in the United States and Germany”. *The Journal of Economic History* 50 (2): 297–307.

Wager, Stefan, and Susan Athey. 2018. “Estimation and Inference of Heterogeneous Treatment Effects using Random Forests”. *Journal of the American Statistical Association* 113 (523): 1228–1242.

Zou, Hui, and Trevor Hastie. 2005. “Regularization and Variable Selection via the Elastic Net”. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67 (2): 301–320.

A. Online Appendix

Contents

A.1. Additional figures and tables

Table A1: Survey response rate

	Any question	All questions asked to all respondents
	(1)	(2)
Treatment	-0.005	-0.004
	(0.002)***	(0.001)**
Control mean	0.226	0.129
N	212277	212277

Notes: This table regresses a dummy equal to 1 if the individual responded to any survey question beyond the first question (column 1) or all questions that were asked to all respondents (column 2) on the treatment dummy. The regression includes agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively.

Table A2: Differences between survey respondents and non-respondents

	Non-respondent mean (1)	Regression coefficient on respondent dummy (2)
Female	48.402	10.040
	[49.975]	(0.265)***
Age		
Under 25 year old	25.097	-9.466
	[43.357]	(0.241)***
Between 25 and 39 year old	46.684	-5.903
	[49.890]	(0.311)***
Between 40 and 54 year old	21.208	9.950
	[40.878]	(0.261)***
Above 55 year old	7.011	5.420
	[25.533]	(0.202)***
Seniority in unemployment		
Less than 3 months	25.807	-2.796
	[43.758]	(0.234)***
Between 3 and 6 months	21.584	-0.937
	[41.140]	(0.237)***
More than 6 months	52.384	3.707
	[49.943]	(0.274)***
Level of education		
No high-school diploma	16.906	-3.273
	[37.481]	(0.194)***
Vocational degree	31.823	-4.180
	[46.579]	(0.249)***
End-of-high-school diploma	19.131	-1.662
	[39.333]	(0.210)***
University degree	32.140	9.113
	[46.702]	(0.259)***
N	164950	212277

Notes: This table reports summary statistics for individuals who did not respond to the survey and the results of tests of equal means with survey respondents for a set of pre-intervention covariates. Column 1 shows mean values for people that did not respond to any question, with standard deviations in brackets. Column 2 reports regression coefficients on the survey response dummy, with standard errors in parentheses. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively.

Table A3: Impact on employment outcomes, for survey respondents

	Survey	6 months after the intervention				18 months after the intervention			
	Self-declared Employment	Employment episode	Stable employment episode	Unemployment duration	Unemployment benefits	Employment episode	Stable employment episode	Unemployment duration	Unemployment benefits
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: ITT estimation									
Treatment	-0.160	-0.319	0.169	0.928	46.907	0.024	0.364	3.038	142.555
	(0.442)	(0.506)	(0.339)	(0.548)*	(56.050)	(0.468)	(0.426)	(1.692)*	(113.708)
Panel B: Instrumental variable estimation: *taker* instrumented with *treatment*									
Taker	-0.388	-0.770	0.409	2.243	113.384	0.058	0.879	7.342	344.587
	(1.067)	(1.222)	(0.820)	(1.318)*	(134.991)	(1.132)	(1.031)	(4.073)*	(273.485)
Control mean	33.679	39.236	14.381	213.183	3936.160	62.131	27.501	470.869	8011.794
N	47327	47327	47327	47327	47327	47327	47327	47327	47327

Notes: Panel A shows ITT results from equation [1] and Panel B shows 2SLS results from equation [2]. All regressions include agency x month fixed effects. Standard errors in parentheses are clustered at the agency level and ***, **, and * indicate significance at 1, 5, and 10% respectively. All regressions except column 1 use dependent variables from administrative data. *Self-declared employment* is a dummy equals to 1 if the job seeker reported that he/she had been employed or is employed at the moment of the survey. *Employment episode* is a dummy equal to 1 for job seekers who experienced at least one employment episode following the intervention. *Stable employment episode* is a dummy equal to 1 for job seekers who experienced at least one stable employment episode (long-term contract, or short-term contract for more than 6 months) following the intervention. *Unemployment duration* counts the total number of days in unemployment since the intervention. *Unemployment benefits* measures the total benefits that job seekers received following the intervention. Regressions in columns 2 to 5 use outcomes measured six months after the intervention. Regressions in columns 6 to 9 use outcomes measured 18 months after the intervention. The sample is restricted to all individuals who responded to any survey question beyond the first one.

Table A4: Impact on reported use of different job search websites and online platforms

	Bob Emploi	Pôle Emploi websites			Private websites				
		Emploi store	La bonne boîte	La bonne formation	Le bon coin	JobiJoba	LinkedIn	CV designer	Jobeggs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: ITT estimation									
Treatment	0.269	0.046	0.024	0.005	-0.002	-0.000	-0.006	0.005	0.001
	(0.005)***	(0.004)***	(0.004)***	(0.003)*	(0.004)	(0.004)	(0.005)	(0.003)	(0.002)
Panel B: Instrumental variable estimation: *taker* instrumented with *treatment*									
Taker	0.645	0.109	0.057	0.013	-0.004	-0.001	-0.015	0.013	0.002
	(0.011)***	(0.010)***	(0.009)***	(0.007)*	(0.010)	(0.010)	(0.011)	(0.008)	(0.006)
Control mean	0.088	0.262	0.130	0.075	0.765	0.234	0.426	0.093	0.058
N	40395	40412	40437	40313	40679	40481	40426	40272	40293

Notes: Panel A shows ITT results from equation [1] and Panel B shows 2SLS results from equation [2]. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. All regressions use dependent variables from survey data.

Table A5: Differences between people using Bob Emploi vs. not using it in the treatment group

	Non-taker mean (1)	Regression coefficient on taker dummy (2)
Female	49.332	9.757
	[49.996]	(0.433)***
Age		
Under 25 year old	24.700	-13.251
	[43.127]	(0.382)***
Between 25 and 39 year old	46.224	-7.689
	[49.857]	(0.479)***
Between 40 and 54 year old	21.817	12.951
	[41.301]	(0.413)***
Above 55 year old	7.259	7.989
	[25.946]	(0.329)***
Seniority in unemployment		
Less than 3 months	25.379	-2.104
	[43.518]	(0.388)***
Between 3 and 6 months	21.736	-2.028
	[41.245]	(0.353)***
More than 6 months	52.662	4.156
	[49.929]	(0.442)***
Level of education		
No high-school diploma	16.574	-3.496
	[37.185]	(0.301)***
Vocational degree	31.775	-4.058
	[46.560]	(0.399)***
End-of-high-school diploma	19.183	-2.340
	[39.374]	(0.330)***
University degree	32.468	9.894
	[46.826]	(0.445)***
N	103821	119525

Notes: This table reports summary statistics for individuals in the treatment group who did not create a Bob Emploi account. It shows the results of tests of equal means with the Bob Emploi users for a set of pre-intervention covariates. Column 1 shows mean values for non-users, with standard deviations in brackets. Column 2 reports regression coefficients on the user dummy, with standard errors in parentheses. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively.

Table A6: Impact on number of applications, 6 months after the intervention

	Number of job applications after 6 months			
	Total	Initiated by job seeker	Initiated by caseworker	Initiated by employer
	(1)	(2)	(3)	(4)
Panel A: ITT estimation				
Treatment	0.023	0.014	0.000	0.002
	(0.066)	(0.064)	(0.007)	(0.002)
Panel B: Instrumental variable estimation: *taker* instrumented with *treatment*				
Taker	0.084	0.053	0.002	0.007
	(0.246)	(0.239)	(0.027)	(0.008)
Control mean	2.439	1.326	0.581	0.089
N	212277	212277	212277	212277

Notes: Panel A shows ITT results from equation [1] and Panel B shows 2SLS results from equation [2]. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. All regressions use dependent variables from administrative data and defined for the period of 8 months after the intervention. The information on who initiated the application is sometimes missing. Therefore, the sum of the ITT effects in columns 2 to 4 differs from the effect in column 1.

Table A7: Impact on number of applications, 18 months after the intervention

	Number of job applications after 18 months			
	Total	Initiated by job seeker	Initiated by caseworker	Initiated by employer
	(1)	(2)	(3)	(4)
Panel A: ITT estimation				
Treatment	-0.026	-0.019	-0.011	0.002
	(0.102)	(0.099)	(0.012)	(0.003)
Panel B: Instrumental variable estimation: *taker* instrumented with *treatment*				
Taker	-0.095	-0.071	-0.039	0.008
	(0.380)	(0.368)	(0.044)	(0.012)
Control mean	4.985	2.930	1.179	0.148
N	212277	212277	212277	212277

Notes: Panel A shows ITT results from equation [1] and Panel B shows 2SLS results from equation [2]. All regressions include agency x month fixed effects. Standard errors are clustered at the agency level, and ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. All regressions use dependent variables from administrative data and defined for the period of 20 months after the intervention. The information on who initiated the application is sometimes missing. Therefore, the sum of the ITT effects in columns 2 to 4 differs from the effect in column 1.

Table A8: Impact on reemployment expectations, robustness checks

	Expects to find a job in less than a month	Expects to find a job in 1 to 3 months	Expects to find a job in 4 to 6 months	Expects to find a job in 7 to 12 months	Expects to find a job in more than a year	Expects to not find a job	Overly optimistic	Overly pessimistic	Realistic
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: ITT estimation without clustering standard errors									
Treatment	0.001	-0.007	0.006	-0.001	-0.003	0.003	0.006	0.000	-0.006
	(0.005)	(0.007)	(0.005)	(0.004)	(0.003)	(0.004)	(0.006)	(0.004)	(0.005)
Control mean	0.185	0.389	0.200	0.098	0.050	0.078	0.693	0.123	0.184
N	23829	23829	23829	23829	23829	23829	23829	23829	23829
Panel B: Westfall-Young multiple hypothesis testing									
SE without clustering		randomized p-val=.862							
SE with clustering		randomized p-val= .862							

Notes: Panel A reports the ITT results from equation [1] but without clustering standard errors. All regressions include agency x month fixed effects. ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. Panel B reports the randomized p-values of the Westfall-Young multiple hypothesis testing based on the 300 randomization iterations we implemented. These p-values are calculated under the null hypothesis that all treatment effects are zero across the outcomes of the table. Overly optimistic job seekers are job seekers that found a job later than what they anticipated. Overly pessimistic job seekers are job seekers that found a job earlier than anticipated. Realistic job seekers are job seekers that found a job within their anticipated timeline.

Table A9: Impact on effort and scope of job search, robustness checks

	Survey data			Administrative data (18 months)			
	Hours per week spent on job search	At least one unsolicited application per week	Applies to jobs outside their municipality	Number of online applications	Number of online applications outside their municipality	Number of online applications outside their target sector	Number of online applications below their reservation wage
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: ITT estimation without clustering standard errors							
Treatment	-0.083	0.004	0.000	-0.019	-0.016	-0.040	-0.005
	(0.066)	(0.005)	(0.005)	(0.092)	(0.085)	(0.057)	(0.046)
Control mean	7.624	0.458	0.505	2.930	2.331	2.112	1.170
N	42985	41855	39117	212277	212277	212277	212277
Panel B: Westfall-Young multiple hypothesis testing							
SE without clustering	randomized p-val=.700						
SE with clustering	randomized p-val= .700						

Notes: Panel A reports the ITT results from equation [1] but without clustering standard errors. All regressions include agency x month fixed effects. ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. Panel B reports the randomized p-values of the Westfall-Young multiple hypothesis testing based on the 300 randomization iterations we implemented. These p-values are calculated under the null hypothesis that all treatment effects are zero across the outcomes of the table. Regressions in columns 1 to 3 use dependent variables from survey data. Regressions in columns 4 to 7 use dependent variables from Pôle emploi's administrative data, which are available for the full sample. Outcomes in these columns are defined for the period of 18 months after the intervention. Outcomes in these columns are defined for the period of six months after the intervention.

Table A10: Impact on job search methods, robustness checks

	Survey data							**Administrative data**	
	Uses personal network	Uses professional network	Adapts CV to job listing	Adapts cover letter to job listing	Follows up with recruting firms	Number of Pôle emploi websites used	Number of private websites used	At least one meeting with caseworker	Number of workshops attended at Pôle emploi
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: ITT estimation without clustering standard errors									
Treatment	0.017	0.008	0.006	-0.003	0.010	0.065	-0.006	0.024	-0.002
	(0.005)***	(0.005)*	(0.005)	(0.005)	(0.005)*	(0.009)***	(0.011)	(0.002)***	(0.009)
Control mean	0.623	0.755	0.608	0.745	0.421	1.369	1.550	0.572	0.653
N	39776	37566	39538	40011	39618	41448	41289	212277	212277
Panel B: Westfall-Young multiple hypothesis testing									
SE without clustering	randomized p-val=.008								
SE with clustering	randomized p-val= .008								

Notes: Panel A reports the ITT results from equation [1] but without clustering standard errors. All regressions include agency x month fixed effects. ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. Panel B reports the randomized p-values of the Westfall-Young multiple hypothesis testing based on the 300 randomization iterations we implemented. These p-values are calculated under the null hypothesis that all treatment effects are zero across the outcomes of the table.

Table A11: Impact on well-being, robustness checks

	Overall well-being	Feels motivated during job search	Feels supported during job search	Partici-pates in non-job-related activities
	(1)	(2)	(3)	(4)
Panel A: ITT estimation without clustering standard errors				
Treatment	0.026	-0.047	0.052	0.003
	(0.022)	(0.023)**	(0.027)*	(0.004)
Control mean	5.167	6.958	3.833	0.772
N	44724	44724	44724	37035
Panel B: Westfall-Young multiple hypothesis testing				
SE without clustering	randomized p-val=.074			
SE with clustering	randomized p-val= .074			

Notes: Panel A reports the ITT results from equation [1] but without clustering standard errors. All regressions include agency x month fixed effects. ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. Panel B reports the randomized p-values of the Westfall-Young multiple hypothesis testing based on the 300 randomization iterations we implemented. These p-values are calculated under the null hypothesis that all treatment effects are zero across the outcomes of the table. All regressions use dependent variables from survey data.

Table A12: Impact on employment outcomes, robustness checks

	6 months after the intervention				**18 months after the intervention**			
	Employment episode	Stable employment episode	Unemployment duration	Unemployment benefits	Employment episode	Stable employment episode	Unemployment duration	Unemployment benefits
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: ITT estimation without clustering standard errors								
Treatment	0.011	-0.025	-0.467	2.642	0.022	0.094	-0.382	12.667
	(0.218)	(0.159)	(0.311)	(19.552)	(0.218)	(0.193)	(0.867)	(40.000)
Control mean	40.485	15.163	199.260	3031.994	59.128	25.578	418.875	6069.893
N	212277	212277	212277	212277	212277	212277	212277	212277
Panel B: Westfall-Young multiple hypothesis testing								
SE without clustering	randomized p-val=.554							
SE with clustering	randomized p-val= .554							

Notes: Panel A reports the ITT results from equation [1] but without clustering standard errors. All regressions include agency x month fixed effects. ***, **, and * indicate significance at the 1, 5, and 10% level, respectively. Panel B reports the randomized p-values of the Westfall-Young multiple hypothesis testing based on the 300 randomization iterations we implemented. These p-values are calculated under the null hypothesis that all treatment effects are zero across the outcomes of the table. All regressions use dependent variables from administrative data. *Employment episode* is a dummy equal to 1 for job seekers who experienced at least one employment episode following the intervention. *Stable employment episode* is a dummy equal to 1 for job seekers who experienced at least one stable employment episode (long-term contract, or short-term contract for more than 6 months) following the intervention. *Unemployment duration* counts the total number of days in unemployment since the intervention. *Unemployment benefits* measures the total benefits that job seekers received following the intervention. Regressions in columns 1 to 4 use outcomes measured six months after the intervention. Regressions in columns 5 to 8 use outcomes measured 18 months after the intervention.

Figure A1: Bob Emploi, example of personalized advice (accessed July 2018)

Figure A2: Bob Emploi, example of diagnosis of employability (accessed July 2018)

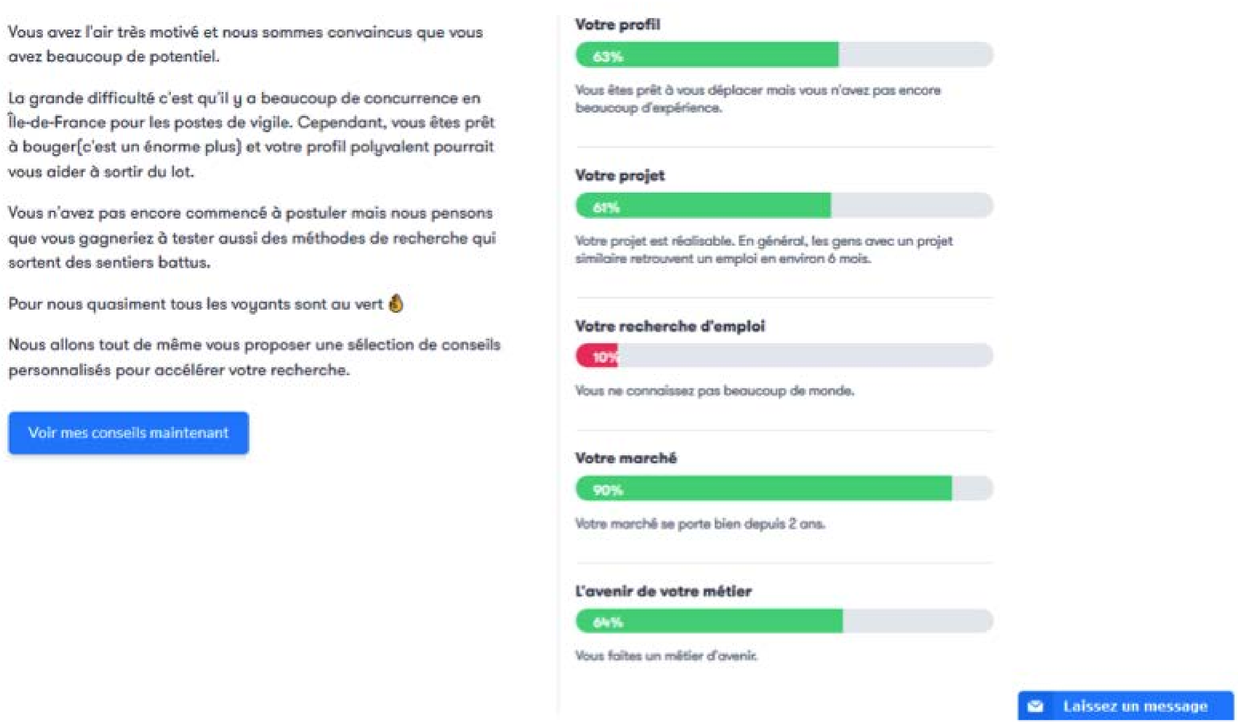

Figure A3: Bob Emploi, example of general advice (accessed July 2018)

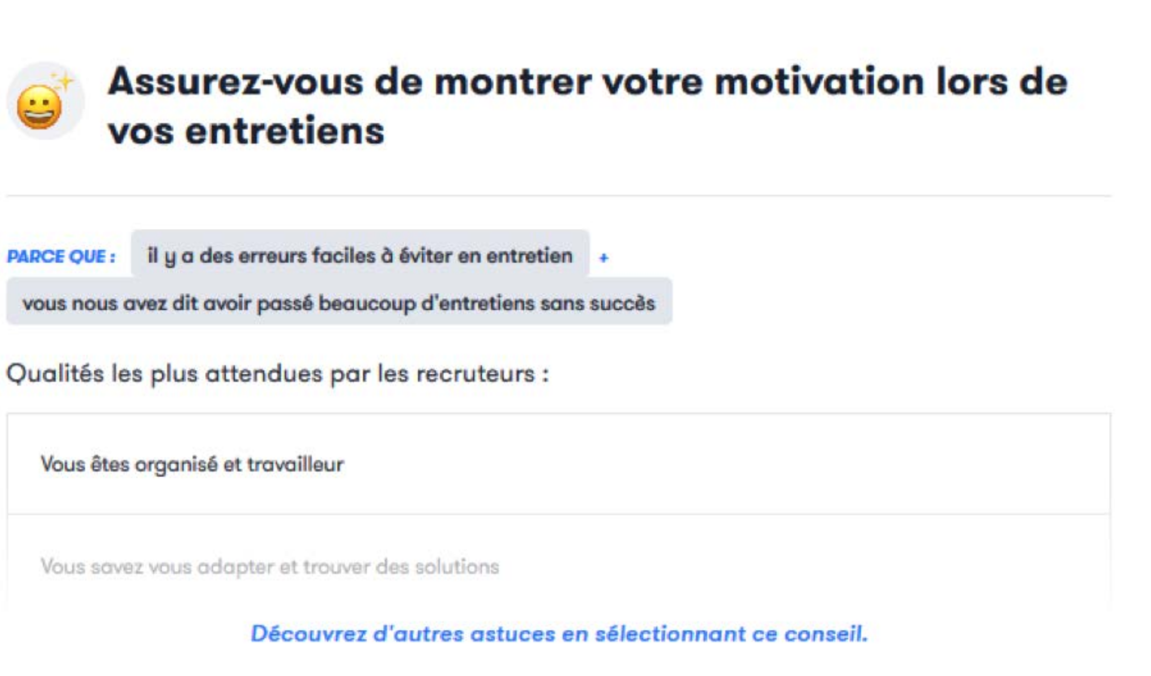

Figure A4: Bob Emploi, example of planning assistance (accessed July 2018)

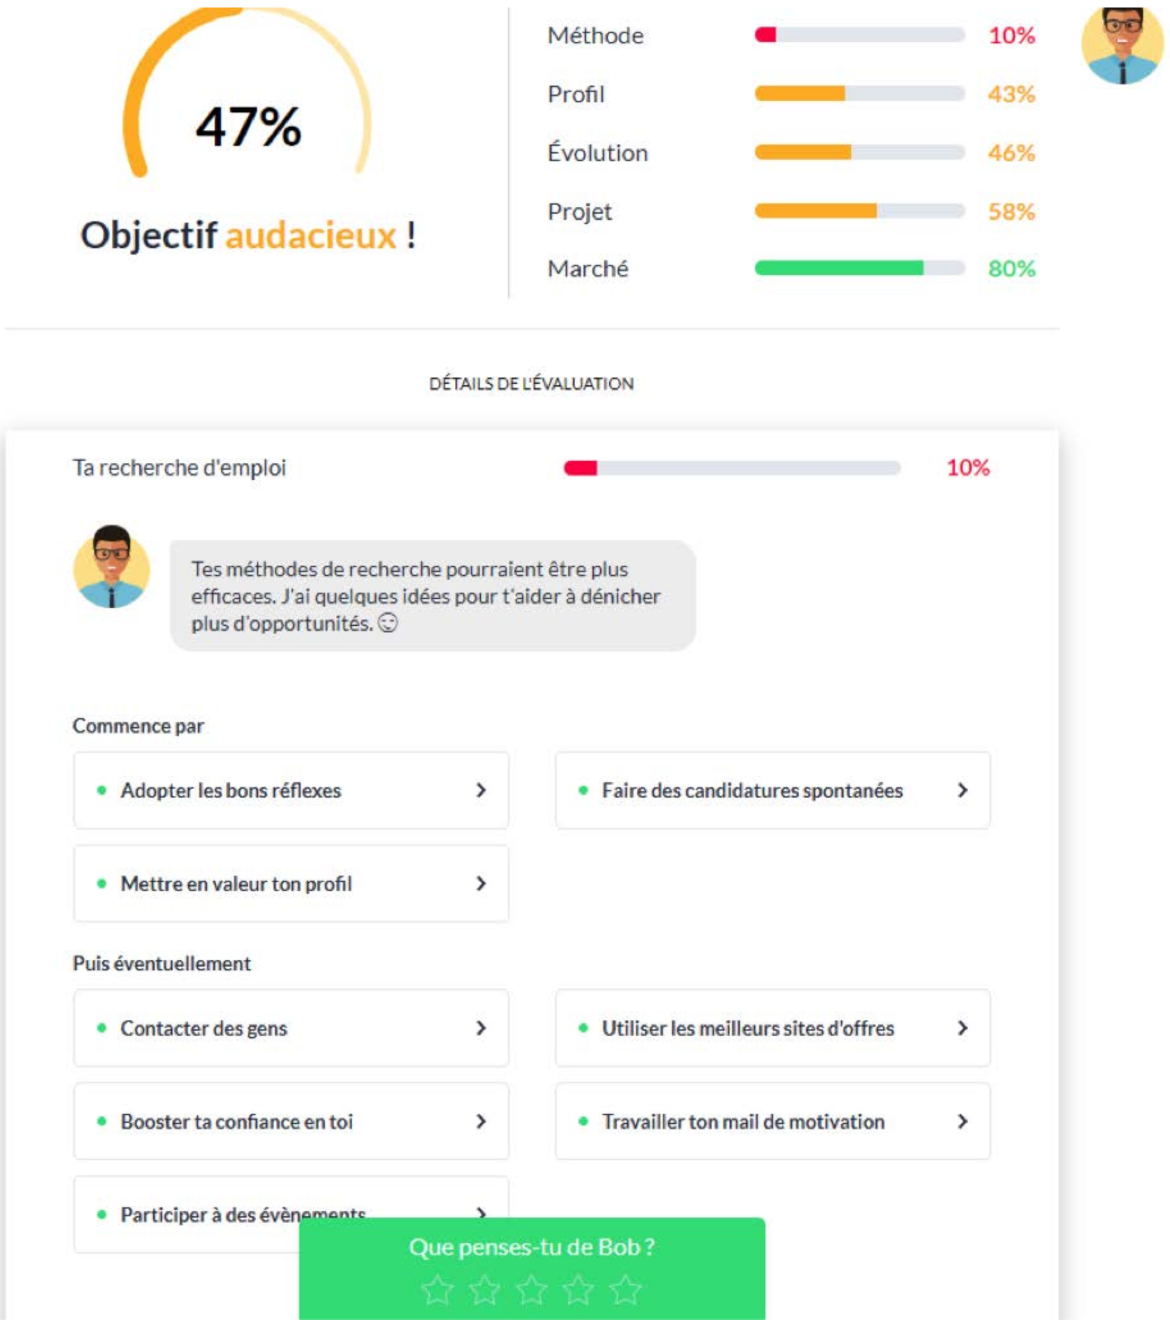

Figure A5: GATES on unemployment duration at six months, generic machine learning

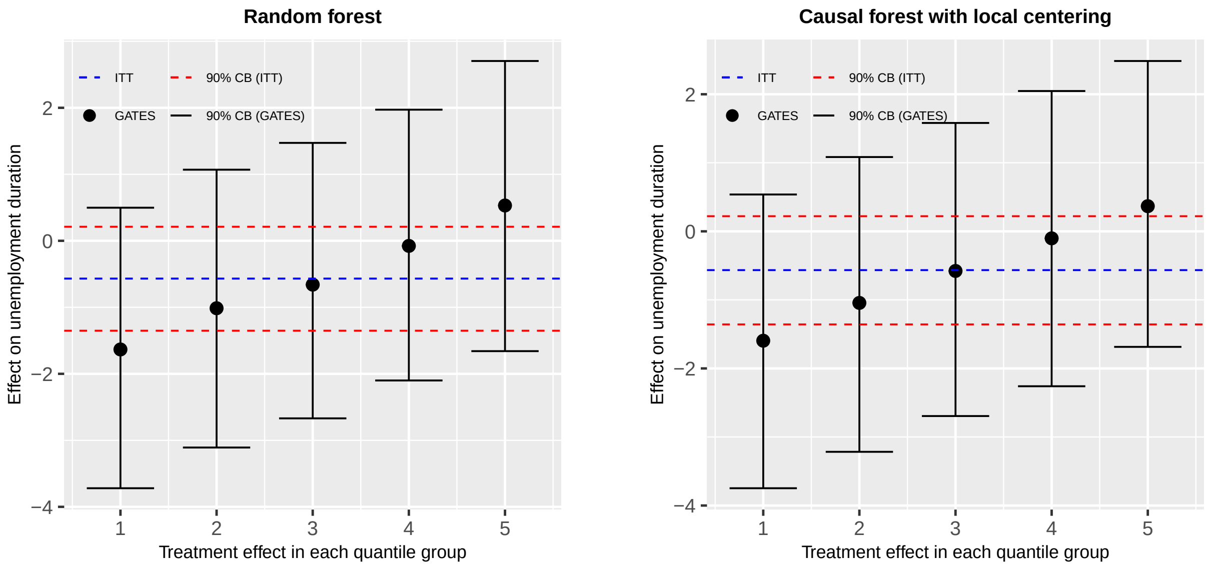

Notes: Sample size = 212,277 individuals. Number of repetitions: 100. We report GATES estimates for random forest (left) and causal forest with local centering (right), which are the two best algorithms to approach heterogeneity according to the selection criteria. Following Chernozhukov et al. (2018), GATES have been reordered after estimation, to make results more readable on the graph. The group heterogeneity score (x-axis) corresponds to an estimated quintile of the ITT. 90% confidence bounds are computed following Chernozhukov et al. (2018). Standard errors are clustered at the agency level.

Figure A6: GATES on unemployment duration at 18 months, generic machine learning

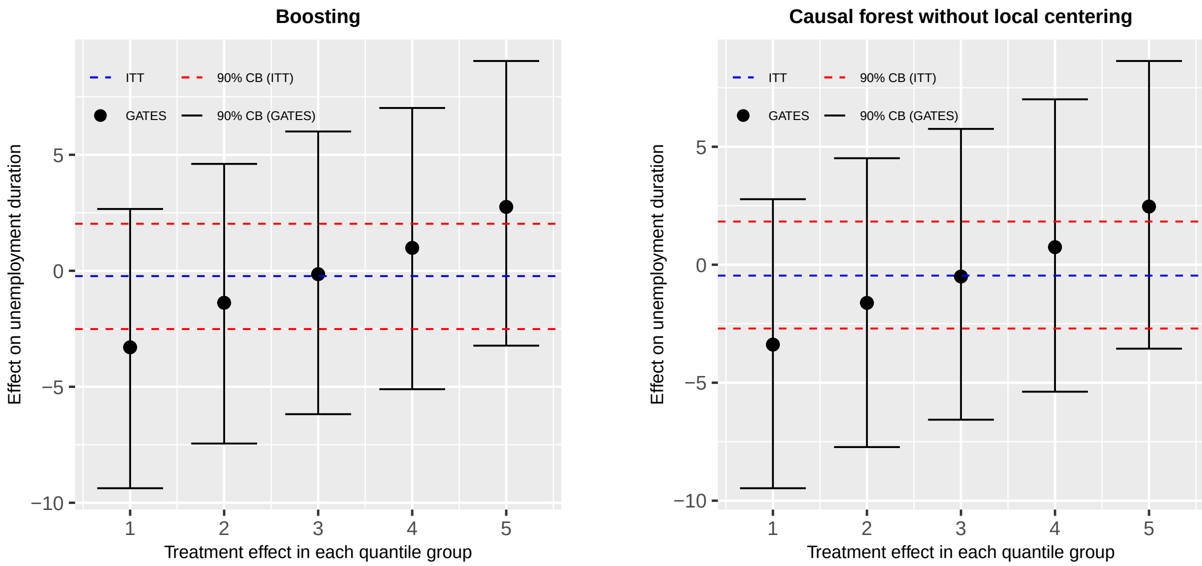

Notes: Sample size = 212,277 individuals. Number of repetitions: 100. We report GATES estimates for boosting (left) and causal forest without local centering (right), which are the two best algorithms to approach heterogeneity according to the selection criteria. Following Chernozhukov et al. (2018), GATES have been reordered after estimation, to make results more readable on the graph. The group heterogeneity score (x-axis) corresponds to an estimated quintile of the ITT. 90% confidence bounds are computed following Chernozhukov et al. (2018). Standard errors are clustered at the agency level.

A.2. Survey questionnaire

The full survey questionnaire is shown below, in its translated version. The first question was directly included in the email that invited job seekers to respond to the survey. All other questions were gathered on a dedicated webpage, accessible via a hyperlink in the email. We do not report the exact format of each question (e.g., whether respondents could fill an empty box or move a cursor to pick a number).

[q1] **What is your current employment status?**

Note: The "employment" category includes all types of contracts (short term / seasonal (CDD), long-term / salaried (CDI), temporary work, and subsidized employment contracts).

- I am unemployed and searching for a job
- I am employed but I am still searching for another job
- I am employed and I am not currently searching for a job
- I intend to create my own business or I am self-employed
- I am doing an internship or a job training
- Other answer (please specify)

This survey takes an average of 5 minutes. Your responses will be kept anonymous and they will never be shared with your caseworker. Your responses are very useful! They help us improve our support services!

The following questions are about your job search.

[q2] **During a typical week, over 7 days, how many days are you active in your job search?**

[q3] **On an average day that you look for work, how much time do you spend on your job search?**

- Between 0 and 30 minutes
- Between 30 minutes and 1 hour
- Between 1h and 2h
- Between 2h and 4h
- More than 4h
- I would prefer not to say

[q4] **In your opinion, what is the most useful action to make progress in your job search?**

- Responding to job offers online

– Networking

– Sending unsolicited applications

– I would prefer not to say

[q5] **Over the past month, how many times have you been invited to interview for a job?**

[q6] **In general, you send unsolicited applications...**

– Multiple times per week

– Once per week

– 1-3 times per month

– Less than once a month

– I have never sent an unsolicited application

– I would prefer not to say

The following 4 questions are about your global outlook on daily life.

[q7] **On a scale of 0-10, how would you rate your life right now?**

Note: 0 indicates that you feel you are living the worst possible life and 10 indicates the best possible life for you. You may slide the cursor to select your response.

[q8] **On a scale of 0-10, how motivated do you feel in your job search?**

Note: 0 represents a total absence of motivation and 10 very strong motivation. You may slide the cursor to select your response.

[q9] **On a scale of 0-10, how much support do you feel you have during your job search?**

Note: 0 indicates a total lack of support and 10 a robust support system. You may slide the cursor to select your response.

[q10] **Please indicate how often you participate in the following activities: multiple times per week / once per week / once per month / multiple times per month / rarely or never / I would prefer not to say**

– A physical activity

– A community activity

– A cultural activity

[q11] **For the following websites, say if: you've never heard of it / you've heard of it / you've used it and find it useful / you've used it and do not find it useful**

– La Bonne Boîte

– L'Emploi Store

– Le Bon Coin

– LinkedIn

– Bob Emploi

– Jobi Joba

– La Bonne Formation

– CV Designer

– Jobeggs

– Pôle-emploi.fr

[q12] **Think back to your latest job applications. How frequently did you do the following? Systematically / Often / Sometimes / Rarely / Never**

– You added key words from the job offer to tailor your résumé

– You modified your cover letter to fit the job offer

– You followed up with the recruiter a few days after sending your application

[q13] **Do you rely on the following people to help you in your job search? Absolutely / More or less / Not really / Not at all**

– Friends and family

– Friends of friends

– Former classmates or college alumni

– Former coworkers

– Your Pôle emploi caseworker

– Volunteers from organizations that assist job seekers

– Local business owners

– People found online (via LinkedIn, etc.)

[q14] **What was your monthly salary (after taxes) in your latest job?**

[q15] **You are searching a job...**

- In your municipality and in neighboring municipalities
- Throughout your département
- Throughout your region
- Throughout the entire country
- I would prefer not to say

[q16] **How much time do you think you will need to find a job?**

- Less than a month.
- Between 1 and 3 months.
- Between 4 and 6 months.
- Between 7 and 12 months.
- More than 12 months.
- I don't think that I will find a job.
- I prefer not to say.

A.3. Implementation details on the machine learning algorithms

Details on the generic machine learning method

Following Rubin (1974)'s framework, let $Y_i(1)$ and $Y_i(0)$ be individual i's unemployment duration if they are in the treatment and control groups, respectively. The impact of Bob Emploi may vary along characteristics X_i. Formally, the Conditional Average Treatment Effect (CATE) can be written as a function of X_i:

$$\text{CATE}(X_i) = \mathbb{E}\big[Y_i(1) - Y_i(0)|X_i\big]$$

Note that the relationship between the characteristics X_i, on one hand, and the individual impact $Y_i(1) - Y_i(0)$ as well as $\text{CATE}(X_i)$, on the other, can be non-linear. In addition, not all characteristics X_i are necessarily relevant to characterize the individual impact, so variable selection techniques may be required. For both reasons, the use of machine learning techniques is a promising avenue to estimate $CATE(X_i)$.

The generic machine learning procedure that we use requires partitioning the data into two datasets of equal size: A and M. The first step is based on using A only and consists in approximating $\text{CATE}(X_i)$ and a second function $B(X_i) = \mathbb{E}\big[Y_i(0)|X_i\big]$. We use two distinct approaches to estimate the function $\text{CATE}(\cdot)$, following the typology of Künzel et al. (2019):

- We use a two-learner and predict separately $\mathbb{E}\big[Y_i(1)|X_i\big]$ in the treatment group and $\mathbb{E}\big[Y_i(0)|X_i\big]$ in the control group. Three classical algorithms are used to perform this type of estimations: gradient boosting, random forest Breiman (2001), and elastic net Zou and Hastie (2005). In this approach, the prediction of $\mathbb{E}\big[Y_i(0)|X_i\big]$ among the control group is also used to approximate $B(X_i)$.

- We directly estimate $\text{CATE}(\cdot)$ with an algorithm designed for causal estimation. The chosen algorithms are the causal forest with and without local centering (Wager and Athey 2018, and Athey, Tibshirani, Wager, et al. 2019), and the Rlearner (Nie and Wager 2020), both with boosting and lasso. In this approach, we estimate $B(X_i)$ separately by restricting A to the control group and using a random forest when we use causal forest, boosting with the Rlearner combined with boosting, and elasting net with the Rlearner combined with lasso. To save computation time, we simply recycle the model that has been fitted in the control group when fitting two-learners.

At the end of the first step, we have approximations corresponding to our functions of interest: $\widehat{\text{CATE}}(\cdot)$ and $\hat{B}(\cdot)$.

The second step of the generic machine learning approach uses the second half of the sample, M. $\overline{\text{CATE}}_M = \frac{1}{\#\{i \in M\}} \sum_{i \in M} \widehat{\text{CATE}}(X_i)$ is the predicted CATE on sample M. Following Chernozhukov et al. (2018), we first run the linear regression:

$$Y_{it} = \theta_0 + \theta_1 \hat{B}(X_i) + \theta_2 \widehat{\mathrm{CATE}}(X_i) + \beta^{\mathrm{ITT}}\big(T_i - p(X_i)\big) + \beta^{\mathrm{HET}}\big(T_i - p(X_i)\big)\Big(\widehat{\mathrm{CATE}}(X_i) - \overline{\mathrm{CATE}}_M\Big) + \mu_{i1}^{a,m} + \epsilon_{it1}, \tag{6}$$

where $p(X_i)$ is the propensity score, $\mu_{i1}^{a,m}$ are strata fixed effects indicating the local agency a x month m of inclusion in the sample, and we use weights $w(X_i) = \frac{1}{p(X_i)(1-p(X_i))}$. Moreover, standard errors are clustered at the local agency level.

Chernozhukov et al. (2018) show that β^{ITT} is the mean treatment effect, and β^{HET} is different from zero if $\widehat{\mathrm{CATE}}(X_i)$ captured heterogeneous effects. Note that $\beta^{\mathrm{HET}} = 0$ can have two possible interpretations: either there is truly no heterogeneous impact, or $\widehat{\mathrm{CATE}}(X_i)$ fits the true $\mathrm{CATE}(X_i)$ too badly to allow us to detect any underlying heterogeneity.[30]

We also run a second regression on sample M, to estimate quintiles of impacts. The general idea is to order the estimated $\widehat{\mathrm{CATE}}(X_i)$ and divide the individuals into five groups of equal sizes: from the lowest approximate impact to the highest.

Specifically, let l_k be the k-th quintile of $\widehat{\mathrm{CATE}}(X_i)$ evaluated on sample M. Let $(Q_{ki})_{k\in[\![1,5]\!]}$ be the binary variables indicating the individual's quintile group, based on the predicted impacts, i.e. $Q_{ki} = \mathbb{1}\{l_{k-1} \leq \widehat{\mathrm{CATE}}(X_i) < l_k\}$, with l_0 simply being the lowest value of $\widehat{\mathrm{CATE}}(X_i)$.

We want to estimate average effects within these quintile groups, which are called Sorted Grouped Average Treatment Effects (GATES). The following linear regression, weighted again by $w(X_i)$, provides unbiased estimates of the GATES γ_k ($k \in [\![1,5]\!]$):

$$Y_i = \theta_0 + \theta_1 \hat{B}(X_i) + \theta_2 \widehat{\mathrm{CATE}}(X_i) + \sum_{k=1}^{5} \gamma_k \cdot (T_i - p(X_i)) \cdot Q_{ki} + \mu_{i1}^{a,m} + \epsilon_{it1} \tag{7}$$

With this approach, not only can we determine whether the impacts within quintile groups are significant, but we can also test for significant differences across quintiles. In particular, we can examine whether the group with the highest predicted impacts is significantly different from the group with the lowest predicted impacts.

Overall, this approach allows us to approximate individualized impacts of the treatment in a very fine and data-driven manner, using machine learning techniques, assess whether there is any detectable heterogeneity thanks to the estimation of β^{HET} and by testing for its significance, and investigate differences in impacts across the distribution of individual impacts.

To reduce the effect of sample splitting on our estimates, we repeat the random partition in A and M, the

[30]We do not comment on θ_0, θ_1, and θ_2, which are coefficients of secondary importance. While the regressor $\widehat{\mathrm{CATE}}(X_i)$ interacted with θ_2 is not explicitly included in the regression shown in the theoretical part of Chernozhukov et al. (2018), it is included in the description of the algorithm and in the authors' code. We follow the implementation of Chernozhukov et al. (2018) and include this control.

training of the algorithms, and the regressions 100 times and we correct p-values and confidence intervals following the guidelines by Chernozhukov et al. (2018).

List of covariates included in the generic machine learning algorithms

- gender - female
- country of citizenship - nationality
- age
- age squared - age_square
- years spent working in the sector of interest - exper_tirage
- whether the job seeker looks for a part-time job or not - temps_part
- whether the individual lives in a sensitive urban area - zus
- reason why the job seeker is unemployed - motif_inscription
- target job sector - famille_prof
- target type of contract (short or long term) - contrat1
- département - departement_de
- availability to take a new job - catregr1
- level of qualification - qualif
- type of support received from Pôle emploi - dc_modalitesuiviaccomp_groups
- distance the job seeker is willing to travel between home and work - distance_quantiles
- the baseline benefits received - tot_amount_groups
- the reservation wage - salmt1_groups
- the number of children - nenf1_groups

Models

Elastic net with expansions

- R Package: glmnet

- R formula:

```
y ~ female * nationality * ( age + age_square + exper_tirage + temps_part + zus + motif_
    inscription + famille_prof + contrat1 + departement_de + catregr1 + qualif + dc_
    modalitesuiviaccomp_groups + distance_quantiles + tot_amount_groups + nenf1_groups +
    salmt1_groups )
```

- Number of individual characteristics: 17 + squared age
- Number of regressors including each interaction: 2 128

- Rules to fit the algorithm using the option `caret`:
  - Method: cross-validation
  - Pre-processing method: range
  - Optimality criteria: default - smallest Root Mean Square Error (RMSE)
  - Number of folds: 2

- Grid of parameters
  - Two parameters to be tuned:
    * α = Mixing percentage between lasso and ridge: 5 values picked using the default grid search methods from the package (`tune = 5`)
    * λ = regularization parameter: 5 values picked using the default grid search methods from the package (`tune = 5`)
  - Grid size: 25 distinct pairs of parameters

Random forest

- R Package: ranger

- R formula:

```
y ~ age + exper_tirage + female + temps_part + zus + nationality + motif_inscription +
    famille_prof + contrat1 + departement_de + catregr1 + qualif + dc_modalitesuiviaccomp_
    groups + distance_quantiles + tot_amount_groups + nenf1_groups + salmt1_groups
```

- Number of individual characteristics: 17

- Rules to fit the algorithm - caret options
  - Method: cross-validation
  - Pre-processing method: range
  - Optimality criteria: default - smallest Root Mean Square Error (RMSE)
  - Number of folds: 2

- Grid of parameters
  - Two parameters to be tuned:
    * Number of randomly selected predictors (`Mtry`): 5 values picked using the default grid search methods from the package (`tune = 5`)
    * Splitting rule (`splitrule`): variance reduction or extratrees
  - Parameters set to default:
    * Minimal node size (`min.node.size`).
  - Grid size: 10 distinct pairs of parameters

Stochastic gradient boosting

- R Package: `gbm`

- R formula:

```
y ~ age + exper_tirage + female + temps_part + zus + nationality + motif_inscription +
    famille_prof + contrat1 + departement_de + catregr1 + qualif + dc_modalitesuiviaccomp_
    groups + distance_quantiles + tot_amount_groups + nenf1_groups + salmt1_groups
```

- Number of individual characteristics: 17

- Rules to fit the algorithm - caret options
  - Method: cross-validation
  - Pre-processing method: range
  - Optimality criteria: default - smallest Root Mean Square Error (RMSE)
  - Number of folds: 2

- Grid of parameters
  - Two parameters to be tuned:
    - Maximum depth of each tree (`interaction.depth`): 5 values picked using the default grid search methods from the package (`tune = 5`)
    - Number of trees (`n.trees`): 5 values picked using the default grid search methods from the package (`tune = 5`)
  - Parameters set to default:
    - Shrinkage parameter applied to each tree in the expansion (`shrinkage`)
    - Minimum number of observations in the terminal nodes (`n.minobsinnode`)
  - Grid size: 25 distinct pairs of parameters

Causal forest with 500 trees

- R Package: grf

- R formula:

```
y ~ age + exper_tirage + female + temps_part + zus + nationality + motif_inscription +
    famille_prof + contrat1 + departement_de + catregr1 + qualif + dc_modalitesuiviaccomp_
    groups + distance_quantiles + tot_amount_groups + nenf1_groups + salmt1_groups
```

- Number of individual characteristics: 17

- Tuning status: default - tuning was attempted but no better parameter than default
- Detailed rules to fit the algorithm can be found here:
  https://grf-labs.github.io/grf/REFERENCE.html#parameter-tuning

---

- Tunable parameters:
  - Tree-growing options:
    - Number of randomly selected predictors (`mtry`): default[31]
    - Fraction of the data used to build each tree `sample.fraction`: 0.5 - default
    - A target for the minimum number of observations in each tree leaf `min.node.size`: default
  - Honesty behavior:
    - Fraction of data used for determining splits to get honest estimates `honesty.fraction`: 0.5 (default)
    - Enable or disable pruning - two possible values `honesty.prune.leaves`: True (default)

[31] Note that when using the grf package, all factor variables are turned to dummies. Thus, mtry can take values 1 to 158, with only 17 different characteristics - including factor variables with several categories.

- Split balance parameters:
  - A tuning parameter that controls the maximum imbalance of a split (`alpha`): 0.05 (default value)
  - A tuning parameter that controls how harshly imbalanced splits are penalized (`imbalance.penalty`): 0 (default value)

Rlearner - lasso

- R Package: rlearner

- R formula:

```
y ~ female * nationality * ( age + age_square + exper_tirage + temps_part + zus + motif_
    inscription + famille_prof + contrat1 + departement_de + catregr1 + qualif + dc_
    modalitesuiviaccomp_groups + distance_quantiles + tot_amount_groups + nenf1_groups +
    salmt1_groups )
```

- Rules to fit the algorithm - caret options
  - Method: cross-validation
  - Pre-processing method: center and scale (default from the `rlearner` package)
  - Optimality criteria: value of `lambda` that gives minimum mean cross-validated error
  - Number of folds: 3

- Number of individual characteristics: 17 + squared age
- Number of regressors including each interaction: 2 128

- Grid of parameters:
  - Propensity score model: automatic generation of `lambdas` - default option
  - Outcome model: automatic generation of `lambdas` - default option
  - Treatment effect model: automatic generation of `lambdas` - default option
  - Note that mixing percentage `alpha` $= 1$: we are fitting lasso models - ℓ^1 penalty only

---

Rlearner - boosting

- R Package: rlearner

- R formula:

```
y ~ age + exper_tirage + female + temps_part + zus + nationality + motif_inscription +
    famille_prof + contrat1 + departement_de + catregr1 + qualif + dc_modalitesuiviaccomp_
    groups + distance_quantiles + tot_amount_groups + nenf1_groups + salmt1_groups
```

- Number of individual characteristics: 17

- Rules to fit the algorithm - caret options
  - Method: cross-validation
  - Pre-processing method: center and scale (default from the `rlearner` package)
  - Optimality criteria: default from `rlearner`
  - Number of folds: 3

- Grid of parameters:
  - Maximum number of trees to grow (`ntrees_max`): 1,000 - default
  - Number of random sampling of hyperparameter combinations for cross validating (`num_search_rounds`): 10 - default

---