Submitted to *Econometrica*

THE POTENTIAL OF RECOMMENDER SYSTEMS FOR DIRECTING JOB SEARCH: A LARGE SCALE EXPERIMENT

LUC BEHAGHEL
Paris School of Economics and INRAE

MARC GURGAND
Paris School of Economics, CNRS, and ENS-PSL

SOFIA DROMUNDO
Organisation for Economic Co-operation and Development

YAGAN HAZARD
Collegio Carlo Alberto and ESOMAS, University of Turin

THOMAS ZUBER
Banque de France

Abstract: We analyze the employment effects of directing job seekers' applications toward establishments likely to recruit. We run a two-sided randomization design involving about 800,000 job seekers and 40,000 establishments, based on an empirical model that recommends each job seeker to firms so as to maximize

Luc Behaghel: luc.behaghel@psemail.eu
Marc Gurgand: marc.gurgand@psemail.eu
Sofia Dromundo: sofia.dromundo@oecd.org
Yagan Hazard: yagan.hazard@carloalberto.org
Thomas Zuber: thomas.zuber@banque-france.fr

We thank Pôle emploi and La Bonne Boîte for their support throughout the project; funding by EUR grant ANR-17-EURE-0001 is acknowledged. We thank Faustin Schneider and Anna Carrere for invaluable research assistance. We thank Philippe Aghion, Axelle Ferrière, Dylan Glover, Thomas Le Barbanchon, Philipp Kircher, David Margolis, Paul Müller, Roland Rathelot, Alexandra Roulet, Gilles Saint-Paul, and numerous seminar participants at Bocconi University, Centre d'analyse et de mathématique sociales (CAMS-EHESS), Collège de France, European University Institute (Florence), Banque de France, INSEAD, AMSE, PSE, Tinbergen Institute, the 2020 "Matching jobs and workers online" IZA workshop, the 2023 EEA meetings, and the 2023 Brucchi Luchino Workshop . All remaining errors are ours. The views expressed in this paper do not reflect the views of the Banque de France. All remaining errors are our own.

total potential employment. Our intervention induces a 1% increase in job finding rates for short term contracts. This impact comes from a targeting effect combining (i) a modest increase in job seekers' applications to the very firms that were recommended to them, and (ii) a high success rate conditional on applying to these firms. Indeed, the success rate of job seekers' applications varies considerably across firms: the efficiency of applications sent to recommended firms is 2.7 times higher than the efficiency of applications to the average firm. This suggests that there can be substantial gains from better targeting job search, leveraging firm-level heterogeneity.

KEYWORDS: Recommender Systems, Matching, RCT, Active Labor Market Policies

.

1. INTRODUCTION

The commercial success of several private recommender systems—Internet-based platforms that go beyond posting job ads or applicant profiles by providing targeted recommendations on potential matches—shows that these services meet a demand on both sides of the labor market, suggesting that they yield positive private returns to firms and job seekers.[1] Can recommendation algorithms be leveraged beyond these private benefits, in order to improve social welfare by reducing search frictions and increasing aggregate employment? Hypothesizing a positive answer, public employment services (PES) have shown increasing interest in providing targeted recommendations, either as an add-on to their main job ads platform, or as separate services. Specifically, based on their profile (if they are logged in) or simply on their actions on the platform, job seekers receive recommendations to expand their search to neighboring occupations or other locations, or to apply to specific firms to which they might not have spontaneously applied.

The rationale is that such services may increase employment by redirecting job seekers both within local labor markets—toward firms with higher hiring potential—and across local labor markets—toward occupations with tighter markets. In these two dimensions, PES can leverage their informational advantage. First, access to past administrative data allows the prediction of the hiring potential of individual firms—which the literature on firm dynamics (Davis et al., 2012, 2013) has documented to be highly heterogeneous. Second, data on the universe of firms and of registered job seekers allow the identification of gaps between the demand and supply of labor at a fine-grained level (e.g., occupation $\times$ commuting zone, our characterization of local labor markets in this paper).

Several conditions are needed for this information to improve labor market outcomes. An obvious one is that job seekers follow the recommendations and redirect their search toward recommended firms and occupations. A second condition is that the recommendations do not crowd out more effective search strategies. In particular, when recommending specific matches, PES lack information on idiosyncratic firms' and job seekers' characteristics that

[1] See in particular Horton (2017), Kuhn and Skuterud (2004), Kuhn and Mansour (2013), Kuhn (2014), Belot et al. (2019, 2022b) for studies of job search platform / recommender systems. Kircher (2020, 2022) provides recent reviews of this literature.

make some matches more productive than others, even within a narrowly defined local market. Further, moving to nearby occupations may generate specific human capital losses that offset the benefits of reduced search frictions. Lastly, as for any active labor market policy, displacement effects remain a concern. If recommendations increase congestion, there is a risk of overshooting by recommending a given firm or local labor market beyond its hiring potential.

In this paper, we provide experimental evidence on the potential value of a large-scale recommendation platform developed by the French PES, combined with a design that optimizes the set of recommendations made to each job seeker, as discussed by Kircher (2022). The platform is called "La Bonne Boîte" ("The adequate firm," henceforth LBB). It was started in 2015, based on an algorithm predicting hirings at the firm $\times$ occupation level. The goal of this service is to leverage data on the universe of firms to identify a subset most likely to hire (the so-called "bonnes boîtes") without necessarily posting job ads at the PES. On its business-as-usual mode, the LBB website directs job seekers toward firms which are predicted to hire and fit their location and occupation criteria (there is no attempt to redirect job seekers toward occupations with tighter markets). We partner with the PES to test the impact of an expanded version of this service using a randomized encouragement design on a pool of 800,000 registered job seekers. We send emails to about 500,000 registered job seekers (the treatment group) to encourage them to use LBB, and recommend them to send applications to specific firms –likely to hire within or outside their occupation– by providing them with links to those firms on the LBB platform. The pool of about 300,000 remaining job seekers forms a control group.

Once we have randomized the treated job seekers in our sample, there can be many ways to match each of them to individual firms. We do so in a way that seeks to maximize job creations. To that end we consider three key factors: (i) firms' heterogeneous hiring dynamics, (ii) occupational switching costs and (iii) firm-level congestion effects. Firstly, we take into account firm-level heterogeneity in hiring dynamics. Job seekers should be recommended more often to firms with better hiring prospects to facilitate job creation in these firms with untapped labor demand. Second, the cost to changing occupation (e.g., human capital loss) generates a trade-off between recommending a firm hiring in a job seeker's origin occupation, versus a firm with a higher hiring potential in a more distant occupation. Lastly, we assume that firm-level labor demand does not respond one-to-one to additional applications. These congestion effects limit the scope of workers redirection toward any given firm. We build a flexible local labor market model that incorporates these three features of the labor market. The model takes into account firm-level predicted hirings, available information on local labor market tightness, and the full distribution of job seekers and firms across occupations. It allows for firm-level congestion effects which depend on the number of applications received by firms and are a source of negative spillover. Solving this model amounts to finding the set of recommendation probabilities of *each* job seeker to *each* LBB firm that maximizes total expected employment. We solve the model on our sample of about 500,000 treated job seekers and 40,000 targeted LBB firms, and draw recommendations based on the optimal recommendation probabilities for each worker-firm pair (employing a so-called Bernoulli trial). This design aims to improve labor market outcomes by increasing applications where they are more effective, both within markets (targeting firms with high hiring potential) and across markets (redirecting job seekers from slack to tight markets). Moreover, the Bernoulli trial provides clean sources of identification to analyze not only the average impact of recommendations, but also which firms it is more effective to target, and whether it is efficient to broaden job search to tighter neighboring occupations.

Using administrative data, we find a positive effect of the e-mails' recommendations and search encouragement on the probability that a job seeker is hired by an LBB firm, and no effect on hiring by other firms. Overall, this implies a 1% increase in job finding rates for short term contracts. To assess whether this is due to a "targeting" effect whereby job seekers are hired precisely in the firms that were recommended to them, we exploit the random variations embedded in our design to assign job search recommendations at the match (job seeker $\times$ firm) level. We find a significant targeting effect: our tailored recommendations increase the likelihood of a given match by 18%, implying that job seekers apply more frequently to the recommended firms, and that these applications are effective.

These effects, though, are limited by the job seekers' application behavior. Leveraging a survey where we asked job seekers if they applied to the recommended firms, we estimate an average application rate of about 7%, and we infer that the intervention only increased it by one percentage point, likely because LBB was already a relatively popular platform, or because these firms were already likely targets. That same information also allows us to estimate the efficiency of sending applications (i.e. the probability that an application sent to a firm results in a recruitment), and its heterogeneity across firms. In that perspective, the *potential* of directing job search is rather high: we estimate that, on average, one in every 143 applications induced by our recommendations would be successful. It is also very heterogeneous: even among recruiting firms identified by LBB, the success rate varies by a factor of 3.4 between firms whose predicted hiring is above vs. below the median; and by a factor of 1.5 between firms belonging to markets with above vs. below median labor market tightness.

This demonstrates that the heterogeneity in firm recruitment behavior, as documented by Davis et al. (2012), can be exploited to improve job seekers' search strategy. Our optimal recommendation model has taken advantage of this heterogeneity, while arbitraging between high hiring potential and occupational distance, and minimizing the potential congestion effects. We find that the average efficiency of applications sent to the firms we specifically recommended is 2.7 times higher than the average efficiency of applications if sent to the population of all LBB firms. This demonstrate the potential of directing search, when recruiting firms can be accurately targeted.

Further, we can analyze occupational switching costs: inspired by the evidence in Kircher (2022), our model tends to recommend jobs that are in the neighborhood of a job seeker's main occupation of search, when they have higher predicted hiring than in the labor market of that main occupation. The randomization design allows us to compare recruitment in such main and neighboring recommended occupations. We find that firms discount the applications of job seekers from neighboring occupations: the success rate of applications to jobs in job seekers' own occupation is 1.5 times higher than in more distant occupations.

Finally, we leverage an additional randomization at the firm level, whereby we varied the number of recommendations assigned to each firm, to explore the congestion generated by directing many applications to the same firm. Unfortunately, as application rates remained low, we did not generate a very strong contrast in terms of applications received. But, although point estimates are imprecise, we find that the success rate of applications to a firm decreases with the number of applications received by that firm, which justifies modeling such congestion when computing a set of optimal recommendations.

This paper contributes to the long-standing endeavor of labor economics to document the benefits of Internet-based job search. Labor economists started paying attention to the potential of the Internet as a match-making device in the early 2000s (Autor, 2001, Kuhn and Skuterud, 2004), with hope but little empirical evidence of its effect on job finding rates. A decade later, further research revived the interest for online job-ads platform with more encouraging observational evidence (Kuhn and Mansour, 2013, Kuhn, 2014). Yet a recent turning point of this literature has lied in the increased capacity to run online controlled experiments to robustly identify and estimate the causal effect of these online platforms on the matching process. Horton (2017) is among the first papers in that strand, highlighting the potential of tailored online screening of applicants to increase the vacancy filling rate on the firm side of the market. A seminal paper by Belot et al. (2019) underscored the potential of personalized online advice to expand the range of occupations in job searches, albeit on a restricted scale. Our work is most related to the very recent and concomitant effort of multiple teams of researchers to partner with PES in several countries (e.g., France, Denmark, Sweden, Switzerland, the UK) to explore at a large scale the potential benefits from online-based assistance to (re-)direct job search (Bied et al., 2023, Altmann et al., 2023, Belot et al., 2022a, 2023, Ben Dhia et al., 2022, Hensvik et al., 2023, Cherubini et al., 2023).[2]

Our paper's contribution is to systematically explore the job creation potential of a simple, low-cost recommender system based on an online platform already used by the PES at scale. In particular, our results underscore the value of making recommendations to job seekers at the firm level—a feature made possible by the use of comprehensive administrative data on past hiring, rather than the sole reliance on job vacancy postings. Indeed, we uncover wide firm-level variations in the propensity to hire, even within local labor mar-

[2]In particular, Belot et al. (2022a) and Belot et al. (2023) indicate that occupational referrals have a positive impact on employment outcomes for long-term unemployed individuals, and for job seekers from structurally slack labor markets. Altmann et al. (2023) also note substantial positive employment effects resulting from information provision regarding job vacancy postings in both job seekers' specific and related occupations—yet they document considerable displacement effects observed when the program is implemented for the majority of job seekers within a specific labor market. Additionally, Hensvik et al. (2023) observe favorable employment effects when directing job seekers' applications toward specific posted vacancies. In contrast, Ben Dhia et al. (2022) find no employment effects from encouraging job seekers to use a private online platform providing tailored job search tips.

kets. This in turn suggests potentially large gains from redirecting job search at the firm level *both* within and across occupations. Such firm-level heterogeneity in hiring dynamics has been previously documented (Davis et al., 2012, 2013), however, to the best of our knowledge, we are the first to explore the extent and manner in which such variation can be leveraged for job search assistance policies so as to ease the matching process on the labor market.[3] Regarding the more frequently studied question of broadening occupational search, our results underscore the risk of overshooting if occupation switching costs are not taken into account—emphasizing the importance of counterbalancing them by targeting substantially tighter labor markets.[4] Lastly, from an applied policy perspective, our flexible matching model provides a workable solution to the complex matching problem when assigning recommendations to job seekers in the presence of congestion externalities and firm and worker heterogeneity. This last point is most related to the work of Bied et al. (2023), who underscore the importance of taking into account competition externalities in the design of recommender systems.[5] Our concomittant work provides empirical evidence on the matter. Moreover, to the best of our knowledge we are the first to implement at scale such a congestion-aware recommender system.

The paper proceeds as follows. Section 2 provides background information on LBB's job search platform as well as a general description of our intervention. Section 3 presents our detailed experimental design as a workable solution to assign recommendations in an ex ante optimal way through emails advertising the platform. Section 4 reports our results both at the individual and job seeker/firm match level. Section 5 concludes.

[3]Recent research by Le Barbanchon et al. (2023) has demonstrated that firm hiring difficulties constitute a significant impediment to firm growth, and generate a relaxation of hiring standards. This supports the hypothesis that guiding job seekers to direct their search efforts towards such firms with unmatched labor demand could ultimately lead to additional job creations.

[4]This aligns to some extent with the recent empirical evaluation of broader job search requirements (van der Klaauw and Vethaak, 2022), that documents potential negative effects of such occupational referrals when they are made mandatory.

[5]Bied et al. (2023) insist on the fact that designing recommendations using state-of-the-art machine learning tools may fail to improve job seekers' outcome if it does not optimize over a *collective* objective function. We derive the optimal recommendation probabilities of our system by maximizing such an aggregate objective function for precisely these reasons.

2. GENERAL CONTEXT AND DATA

2.1. *"La Bonne Boîte," an online job search platform*

This study builds upon a pre-existing platform, "La Bonne Boîte" (LBB). This platform has been operated by the French Public Employment Service (PES) since 2015, that is for five years before the experiment presented in this paper. In this section, we briefly review the main pre-existing features of the platform.

LBB is an online job search platform that aims to help users in their search by directing them toward firms with a high hiring potential. It is presented as a tool to make effective unsolicited (spontaneous) applications (the site is marketed on Google with the motto "Don't send your résumés randomly anymore!"). It can be accessed by any job seeker without registration, and works as a search engine: job seekers indicate a geographical area and an occupation of search (see Figure A1) and LBB proposes a list of firms likely to hire them (see Figure A2). Once they click on a firm of interest, an email address and/or phone contact of the firm is provided (see Figure A3).

The distinguishing feature of LBB is to recommend firms deemed likely to hire, whether they have posted a job vacancy or not. To do so, LBB uses administrative data covering the universe of French firms to derive hiring predictions at the establishment $\times$ occupation level.[6] LBB then defines for each occupation a specific predicted hiring threshold above which an establishment is deemed a "hiring firm" for this specific occupation.[7] If there is no such establishment, LBB's search engine suggests to extend the search to a wider geographical area. We do not have leeway on the algorithm used to predict hiring, and take it as given. However, we check that the quality of LBB's prediction is sufficient for our purposes.[8]

[6]These predictions are derived from establishment level predictions which are then mapped into establishment $\times$ occupation hiring prediction using a sector-occupation crosswalk. This crosswalk is based on the share of each occupation hirings within each sector. This share was computed for registered unemployed exiting unemployment between the 02.03.2016 and 31.03.2017 (https://www.data.gouv.fr/fr/datasets/nombre-dembauches-par-code-ape-et-code-rome/).

[7]As a consequence, a given establishment can be considered as a "hiring firm" for one occupation but not for another.

[8]Figure A6 in the Appendix plots the relationship between the log of firms' average predicted hiring as of August 2019, within twenty equal-size groups, and the log of realized average hiring in each of those groups of

2.2. *Emailing job seekers with tailored recommendations*

In practice, our experiment consisted in emailing treated job seekers with links to a small number of firms on the LBB platform, randomly selected among the firms that fit each job seeker's geographical location and her occupation of search or a neighboring one. Job seekers interested in the recommended firms were encouraged to contact them to make an application. The contact information usually consists of a location, an email or a telephone number. Moreover, in some cases LBB allows job seekers to directly send an application through the PES online application tool. When this tool is available, job seekers simply need to click on the "Send an application" (in French "Postuler") icon, as can be seen in Figure A3 in Appendix A.1.

The experiment took place between November 19th 2019 and December 4th 2019. During this period we sent more than $2,400,000$ emails to the pool of treated job seekers, in four different batches. As can be seen in Table I below or in Figure A4 in Appendix A.1, the emails contained the following information: the job seeker's name, the statement that a considerable share of hirings stem from unsolicited applications (so as to encourage job seekers to apply to these firms even in the absence of a posted vacancy),[9] general information on LBB, each job seeker's declared occupation of search, at most two links to the LBB page of recommended firms and, finally, a general purpose link directing toward LBB's search engine.

Using a design that will be detailed below, we drew up to eight firms within the pool of LBB firms to recommend to job seekers by email. Because we were unsure about how many firms we should recommend to a job seeker, we randomly drew job seekers to receive either two or four emails, with at most two different recommended firms in each email.[10] Finally, we distinguished between firms hiring in a job seeker's own occupation, and firms hiring in

firms during the six following months. The figure also plots the linear correlation between the logs of predicted hiring and realized hiring, estimated on the individual data. The correlation coefficient is 0.89, with an R-squared of 0.37, and significant at the 1% level.

[9]Recall that the selection of firms on the LBB platform, thus also our recommendations, are based on predicted hiring behavior and do not use any information about posted vacancies.

[10]We did not find a statistically significant difference between the outcomes of job seekers receiving two or four emails. Firms to be recommended were drawn independently; when a single firm was drawn to appear twice in a single email, we collapsed the two links into one single link.

a neighboring occupation by introducing the links with a different framing: establishments hiring in one's own occupation were introduced as such, whereas establishments hiring in a neighboring occupation were framed as "hiring a profile close to yours" (in French "Un profil proche du vôtre").

TABLE I
EMAIL'S STRUCTURE

Dear Mr./Mrs. [X], You are currently registered with the public employment service and are looking for a job as a [X's occupation of search]. Did you know that 7 out of 10 firms take into consideration unsolicited applications before actually posting a job-offer? "La Bonne Boîte", an online platform linked to the Public Employment Service, has selected for you several firms which might be interested in your profile. Here is one that is likely to be interested in [your profile/a profile close to yours]: - [Link to recommended establishment 1] And another one that is likely to be interested in [your profile/a profile close to yours]: - [Link to recommended establishment 2, if any] You can send them your application. By clicking on [this link/these links] you will be able to contact [this firm/these firms] thanks to the coordinates that will appear or by using PES' online application tool if it is available. You may also search for other firms on LBB's website [general purpose link] Yours sincerely,

2.3. *Data*

Firms. On the firm side, we use LBB's data, which include the number of predicted hirings per occupation and establishment, an indicator of the fact that the firm is identified as a "hiring firm", the firm national identifier, and its location (ZIP Code). Our initial sample consists of $98,366$ LBB hiring firms. We select at random a subset of $38,810$ of these LBB firms which we use to make tailored job search recommendations in the experiment.[11]

[11] We do not insist here on the firm-level randomization, whose analysis is the focus of a companion paper. It is sufficient to mention that we stratify the random selection of firms within 5-digit sectors and above median/below

Job seekers. We exploit exhaustive administrative data from the PES. Besides important demographic characteristics (gender, age, experience, diploma, nationality, etc), and information on the past and current unemployment spells, we know from where and in which specific occupation (ROME code) each job seeker is currently searching. This occupation may or may not be identical to their previous occupation. This data source also provides the main outcome of interest: hirings (date and type of contract, and employing firm) obtained through employment declarations that employers are legally bound to fill ("DPAE").

After dropping all job seekers whose desired occupation is missing ($274,662$), all job seekers for whom we were unable to get a valid email address ($198,510$) and all job seekers listed as currently unavailable for active work ($609,547$), we obtain a final sample of $1,209,859$ active and registered unemployed job seekers. We accessed the data on September 30^{th}, 2019 and dropped ex post all job seekers who had left unemployment between that date and the start of the experiment (November 19^{th}, 2019).

Survey. To get some insights on job seekers' reactions to the emailing campaign, we ran a short web survey in a representative sample of 11,741 job seekers. Outcomes are measured about two months after the emails were sent. We asked job seekers about their usage of the LBB platform, responses to job ads, the number of applications (unsolicited or not). As is common with such web surveys, the response rate was relatively low (26%). We account for sampling and non-response weights in all measures taken from this survey.

Descriptive statistics. Table II describes the job seeker samples. Based on the administrative data, 45% are male, 61% hold a high school diploma, the average age is about 39, the average work experience 6.9 years and the average unemployment spell at the time of the experiment is 21 months. Panel B of Table II describes the search behavior of control job seekers (in the absence of email encouragement to use the LBB platform and send applications to a set of recommended firms). A vast majority is already using several Internet search channels, including the LBB platform for 20% of them. About half of job seekers report having made unsolicited applications over the previous two months—a proportion

median predicted hiring bins. The heterogeneity analysis in Section 4.5 is made along the latter stratification variable.

similar to the share having responded to job ads—and the unconditional average number of applications are of the same order of magnitude (3.86 unsolicited applications vs. 4.79 responses to job ads). This shows that unsolicited applications are already part of the job seekers' strategies before our intervention—perhaps partly due to the already relatively widespread use of the LBB platform.

TABLE II
JOB SEEKER SAMPLE DESCRIPTION

	(1)	(2)
	A. Job seeker characteristics	
Male	0.45	(0.5)
Age	38.94	(12.05)
Graduated from high school	0.61	(0.49)
Years of experience	6.92	(8.2)
Unemployment duration (months)	21.26	(24.72)
Predicted exit rate from unemployment	0.21	(0.07)
Number of observations	800,297	
	B. Job search	
Used Internet for job search	0.86	(0.35)
# Internet search channels used	2.46	(1.51)
Used LBB	0.20	(0.4)
Responded to job ads	0.54	(0.5)
# job ads responded	4.79	(11.36)
Made spontaneous application	0.51	(0.5)
# spontaneous applications	3.86	(8.33)
Applied in other occupation	0.49	(0.5)
# hours searched per week	8.38	(11.21)
Number of observations	1,102	

Note: Column (1) displays sample average, column (2) displays standard deviations. Source: Administrative data (panel A) and job seekers' online survey (panel B). In panel B, only control job seekers are included.

3. EXPERIMENTAL DESIGN

Our experiment consisted in emailing treated job seekers with links to the contact information of a subset of firms listed on LBB. Interested job seekers could then use this contact information to send applications. The design of the experiment implies that there are two main levels of randomness in our treatment. First and foremost, some job seekers received our recommendation emails (the treated group) and some others did not (the control group). This level of randomness allows us to identify the overall effect of our intervention on job seekers' subsequent labor market outcomes. The second level of randomness is that of the actual recommendations. Treated job seekers were recommended specific firms in a conditionally random manner. This second level of randomness—which exact set of firms gets recommended to a given job seeker—allows us to identify the effect of a targeted redirection of job seekers' search effort. This section describes these two different levels of treatment.

3.1. *Drawing treated job seekers*

Our experimental sample covers 94 randomly selected commuting zones out the of the 404 French commuting zones, representing a pool of $1,209,859$ job seekers. We randomly assign two thirds of the job seekers within these 94 commuting zones to be treated, i.e. receive the emails pushing the LBB service, with specific recommendations toward LBB firms.[12] We stratify the random selection of treated job seekers within commuting zones, reported occupation of search and above median/below median bins of a linearly predicted exit rate out of unemployment using predetermined worker information.[13] We randomly allocate $806,437$ job seekers to the treatment. Because a large share of job seekers exited

[12]When assigning treatment within a commuting zone, we do not take into account the geographical distance between job seeker and establishment pairs. Indeed, the existing evidence suggests that spatial mismatch is of second order compared to occupational mismatch (Marinescu and Rathelot, 2018). We stratified the random selection of treated commuting zones within labor market tightness and size quintiles. For more details on commuting zones and local labor markets, see Appendix Section A.4.

[13]We predict the exit rate out of unemployment within six month trough a simple linear probability model on job seekers' observable variables (gender, age, level of education, qualification etc.) in an historic version of our administrative data set which encompasses the job finding history of all registered unemployed job seekers between 2016 and 2018. We use the predictions of this model in our sample as a synthetic index. This allows us to reduce the number of stratification variables while still improving the balance between control and treatment group.

the unemployment pool in the short period between randomization and the actual start of our experiment, we restrict our analysis ex post to the $533,557$ treated and $266,740$ control job seekers who were still registered with the PES and had not found a job as of November 19th, 2019.[14] The balance of job seekers' observable predetermined variables across treatment and control groups is presented in Appendix A.2, Table A1. Treated and control job seekers are shown to be similar along a wide set of observable dimensions.

3.2. *Matching job seekers and firms*

There are many ways to match job seekers and firms. Within commuting zones, some firms are predicted to hire more than others. What's more, from the point of view of a given job seeker, firms are predicted to hire in occupations which are more or less close to his or her own preferred occupation.[15] In order to reduce the degree of labor market mismatch one would ideally need to trade-off the adjustment cost of switching occupations with the gains associated to a strong labor demand coming from firms in neighboring occupations, but in tighter markets.

Matching model To solve this trade-off in practice we build a simple matching model as a means to generate sensible pairwise recommendation probabilities in a principled way. This model describes the outcome of job seekers' applications conditional on our recommendations. More specifically we assume that job seekers' probability to apply to a given firm as well as firms' probability to hire a given job seeker are both decreasing functions of occupational distance. We further assume that firm-level congestion effects arise out of a concave application screening technology. In practice, this last feature of our model prevents us from over-flooding firms with too many job applications. The model is solved to

[14] This pre-treatment attrition rate is well balanced across treatment and control groups. The high attrition is due to the delay it takes to the PES to consolidate the database with job seekers' characteristics that we used in our stratified randomization.

[15] We measure the distance between any two occupations as the shortest path in the network of occupations defined by the set of "close" occupations according to Public Employment Services. "Close" occupations are occupations between which job seekers are able to transition without any form of retraining. Linking close occupations together we construct the network of occupations implied by skill proximity in the French ROME classification (532 occupations). The resulting measure of occupational distance is discrete, ranges from $d = 0$ to $d = 19$ in our sample and is well correlated with other measures based on skill classifications such as O*Net.

maximize the total expected employment of job seekers for reasonable values of the key parameters (cost of occupational mobility, concavity of the firm screening function), separately in each commuting zone. The specification of the model and its parameterization are detailed in Appendix A.6.

Formally, we characterize the set of recommendation probabilities $\pi_{i,j} = P(R_i^j = 1)$ that maximize the expected number of job creations in each commuting zone according to our model, and given the empirical structure of supply and demand in all occupations in the zone, where R_i^j is equal to 1 when we recommend firm j to job seeker i and 0 otherwise. Those probabilities are conditional on the firm's predicted hiring and the job seeker's occupational distance to that firm. Indeed, the model penalizes the recommendation of firms further away in the occupational space, whereas it rewards the recommendation of firms with higher predicted hirings. This can be seen in Figure 1 showing that the resulting recommendations probabilities $\pi_{i,j}$ are decreasing in our measure of occupational distance (Figure 1a) and increasing in the level of firms' predicted hirings (Figure 1b). This design solves the multidimensional problem of allocating firms to workers in a way that is policy relevant.

To sum up, our matching model makes a principled compromise between two polar alternatives: the business-as-usual mode of the LBB platform that restricts recommendations to firms and job seekers within the same occupation, on the one hand, and fully random recommendations made irrespective of occupational distance, on the other hand. While these two extreme alternatives greatly simplify the matching problem, the first one may end up increasing congestion (by multiplying search effort where job seekers are already in excess supply), while the second may end up increasing search costs (by recommending unrealistic matches, given the occupation switching cost that they would imply).

FIGURE 1.—CORRELATIONS OF RECOMMENDATION PROBABILITIES π_{ij} WITH OCCUPATIONAL DISTANCE AND WITH PREDICTED HIRINGS

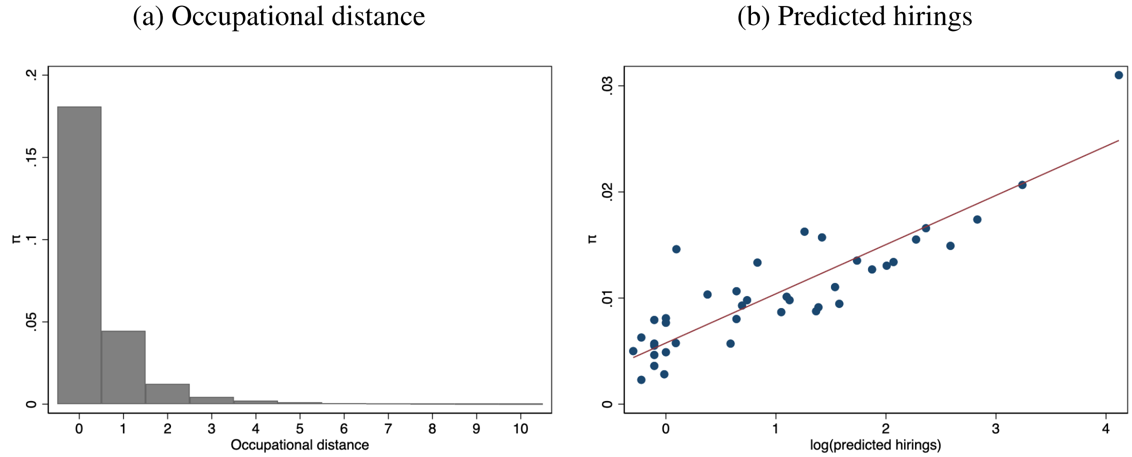

Notes: Panel (a) displays mean recommendation probabilities by occupational distance, panel (b) displays a binned scatter plot of recommendation probabilities on the log of establishment level predicted hirings. In panel (a) recommendations probabilities sum to 1 within job seeker across firms and occupations, but not across occupational distances. To obtain the distribution of recommendations by occupational distance, the mean probabilities reported in panel (a) should be re-scaled by the number of potential recommendations at a given occupational distance divided by the total number of recommendations actually made. Overall, 61.7% of our recommendations were drawn in job seekers' preferred occupations ($d = 0$), 15.6% in occupations listed as "close" by the PES ($d = 1$) and $14,4\%$ in more distant occupations ($d > 1$).

Assignment of recommendations With this set of recommendation probabilities in hand we then draw recommendations $R_i^j \in \{0,1\}$ independently for each pair (i,j). By construction of our experimental design the recommendation dummy $R_{i,j}$ is orthogonal to hiring counterfactuals conditional on the propensity score $\pi_{i,j}$, allowing causal quantities to be identified by propensity score re-weighting (Rosenbaum and Rubin, 1983).

Finally, in order to later assess the degree to which firm-level congestion effects might limit the impact of tailored job search recommendations outside of the equilibrium generated by this design, we randomly vary the number of times that a given firm gets included in our recommendation emails. Concretely, we allocate recommended firms into two different treatment arms, labeled "few" and "many" respectively. While we randomly attribute to the first group of firms a low job applications screening efficiency parameter, we endow the second group of firms with a high value of this same screening efficiency parameter. As a consequence, the optimal recommendation probabilities $\pi_{i,j}$ for firms in the "few" recommendations group are approximately half as large as those of firms randomly allocated to the "many" recommendations group, because the latter face a lower cost of screening many applications when we solve the model.[16] We use this random variation to identify the congestion effects which could potentially affect firms as they are recommended more job seekers, thus likely to receive more applications (Section 4.6).

4. RESULTS

In this section, we first show that our recommendation emails were received, read and used by a significant number of job seekers. We then provide reduced-form evidence that receiving the email increased job finding rates in LBB firms. Finally we leverage our match level randomization design to disentangle the possible mechanisms at play in our experiment. This allows us to distinguish and identify the respective roles played by the "targeting" of recommended firms by treated job seekers as opposed to other indirect effects of our intervention. We show that this targeting effect is large on average (it increases by 18% the probability that a match occurs), with significant variations due to differences in application success rates depending on the firms' hiring forecast, and, to a lower extent, on the occupational distance between the firm and the job seeker, and on the local labor market tightness.

4.1. *Take-up of the treatment*

The email sender keeps track of emails received, opened, and if the links were clicked. Using the survey we can also judge in what proportion the job seekers applied to the recommended firms. This is summarized in Table III. With the tracking data, we observe that job seekers clicked on the proposed links 25% of the time, and often on several links. In the survey, we ask treated job seekers whether they contacted the firms recommended to them in the emails they had received from LBB. On average, treated job seekers report that they

[16] The average recommendation probabilities are respectively 0.006 and 0.012 in the "few" and "many" treatment arms. See Appendix A.6 for more details on the implementation of this random variation through our labor market model.

have contacted 7.3% of the recommended firms.[17] Conditional on having clicked at least one link, applications were sent to 13% of the recommended firms; but even job seekers who did not click did send applications to 6% of them. Of course, they also could have applied to those firms even in the absence of recommendations: we discuss in Section 4.3 how this counterfactual can be recovered. Further, our survey indicates that about 20% of job seekers know and use the LBB platform: this increases by about 3 percentage points in the treatment group, and may imply, as a side effect, that treated job seekers search more, on the LBB platform or even generally.

[17]Job seekers were recommended several firms. We list up to five recommended firms, and ask separately for each of them whether the job seeker has contacted it: 12.1% report that they have contacted exactly one firm, 4.3% report having contacted two firms, and 0.7% report having contacted three or more firms.

TABLE III

TAKE-UP OF THE TREATMENT

	Mean	Sd.	N
A. Tracking data			
(i) Received email	0.96	0.19	533,557
(ii) Opened email	0.64	0.48	533,557
(iii) Click	0.25	0.43	533,557
(iv) Click if opened email	0.36	0.48	340,777
(v) Total clicks if click	2.98	3.02	130,810
(vi) Distinct clicks if click	1.95	1.09	130,810
B. Job seeker survey			
(vii) Application rate	0.073	0.260	8,061

Notes: The first three lines of the table report the rates at which treated job seekers (i) received the e-mail sent on a well-functionning e-mail address (96%), (ii) opened an e-mail (64%), and (iii) clicked on any of the link contained in our e-mails (25%). These are all unconditional rates, over the whole population of treated job seekers in our experiment. The next three lines display (iv) the rate at which job seekers clicked on our links conditional on opening one of our e-mails (36%), (v) the average number of clicks on any of our links conditional on clicking once (2.98), and (vi) the average number of clicks on distinct links conditional on clicking once. Lastly, line (vii) reports the share of recommended firms to which treated job seekers have applied, based on the job seeker survey.

4.2. *Intention-To-Treat impacts on job finding rates*

We observe access to employment over a period of four months since treatment (until the start of the Covid-19 pandemic),[18] for all job seekers in our treatment and control group. We also observe the characteristics of the contract and the firm hiring them, such as whether they are hired (i) in a short-term or long-term contract, or (ii) in a LBB firm (i.e., a firm displayed on LBB platform based on its high predicted hiring score) or non-LBB firm. The effect of our intervention on the job finding rate of job seekers is identified by the mere comparison of control and treated job seekers' outcomes. Formally, we denote by $Y_i \in \{0,1\}$ the indicator variable equal to 1 if job seeker i has been hired during the four months following the launch of our intervention, and $Z_i \in \{0,1\}$ the indicator equal to 1 if i belongs to the treatment group. Following the usual potential outcome framework, we also define $Y_i(z_i = 1)$ and $Y_i(z_i = 0)$ being (respectively) the potential outcomes of job seeker i if treated or not. The Intention-To-Treat (ITT) effect of our intervention, is thus defined as:

$$\mathrm{ITT} \equiv \mathrm{E}\left[Y_i(z_i = 1) - Y_i(z_i = 0)\right].$$

As our randomization ensures that Z_i is independent from the potential outcomes $Y_i(z_i = 1)$ and $Y_i(z_i = 0)$, the ITT parameter is straightforwardly identified by the difference in mean outcomes between treated and control job seekers.

The estimates of the ITT on job finding rates—further disaggregated by type of contract—are reported in Table IV. We do not observe a significant impact on the overall job finding rate (column (1)). However, there is a marginally significant (p-value = 0.09) impact of our intervention on job finding rates in short-term contracts, the most frequent outcome at baseline (column (3)). The magnitude of this impact (0.14 percentage point) is economically meaningful, representing a 1% increase in the baseline job finding rate in short term contracts.

Beyond these average effects, systematic investigation of heterogeneity along job seekers' observables does not uncover significant patterns. In particular, we implement two

[18] We restrict our attention to this horizon as the lockdown induced by the Covid-19 pandemic started on the March 13th 2020 in France, massively disturbing labor market dynamics.

TABLE IV

EFFECT ON JOB FINDING BY CONTRACT TYPE (ITT)

	(1) All	(2) Long term	(3) Short term
Treated (Z_i)	0.0008	-0.0007	0.0014
	(0.0009)	(0.0005)	(0.0008)
	[0.42]	[0.16]	[0.09]
Baseline	0.19	0.04	0.15
Observations	800,297	800,297	800,297

Notes: Standard errors clustered at the labor market level (CZ × Occ.) reported in parenthesis.

agnostic ML-based tests for treatment effect heterogeneity (Chernozhukov et al., 2018, Yadlowsky et al., 2021) that do not conclude to significant levels of heterogeneity. More classical explorations do not conclude to large heterogeneity either.[19] All these results are reported in Appendix A.7.

4.3. *The targeting effect*

Unboxing possible mechanisms There are four mechanisms by which the intervention may affect job-finding. First, job seekers may have sent applications to the very firms that we recommended them in the email. Second, in that case, they could have also substituted the applications they would have made with the recommended ones. Third, they may have used the LBB platform more intensively, and successfully applied to other firms on this platform (given that the emails encouraged the use of the LBB platform, beyond providing recommendations). Fourth, the email may have prompted their search effort in general, even outside LBB. Table V seems to exclude the latter interpretation. It splits the ITT estimate on short term contracts from Table IV column (3) into two parts: an effect on short term job finding into firms absent from the LBB platform and an effect on short term job finding into firms present on the LBB platform.

[19] If anything, we find limited evidence of larger employment effects (in short term contracts) for women, more educated, and long-term unemployed individuals.

TABLE V

EFFECT ON SHORT TERM JOB FINDING BY TYPE OF FIRM (ITT)

	(1) All	(2) Not LBB	(3) LBB
Treated (Z_i)	0.00142	0.00030	0.00112
	(0.0008)	(0.0007)	(0.0005)
	[0.09]	[0.67]	[0.04]
Baseline	0.154	0.097	0.057
Observations	800,297	800,297	800,297

Notes: Standard errors clustered at the labor market level (CZ $\times$ Occ.) reported in parenthesis. The impacts on job finding rate in short term contract displayed in the first column are decomposed into different categories of hiring, depending on the type of recuiting firm. Column 2 (Not LBB) reports hires in short term contract from non-LBB firms—i.e., firms without high enough predicted hirings to meet the bar of LBB's algorithm. Column 3 (LBB) focuses on LBB firms —i.e., firms with high predicted hirings from LBB's algorithm.

The small average effect that we find on overall short term job finding (about 1% of the baseline) comes from increased short term job finding in LBB firms present on the platform. For this group of firms, the effect on short term job finding is close to 2% of the baseline. In contrast, the effect is small and not statistically significant for firms which were not included in the platform.

However, the positive effect on LBB firms does not disentangle the first three mechanisms listed above. Also, under the assumption that LBB firms were more likely to hire than non-LBB firms, even an overall increase in treated jobs seekers' search effort could well have resulted in a positive differential job finding rate only concentrated on LBB firms. In order to disentangle this, we go beyond worker-level job finding effects and estimate pairwise worker/firm level effects.

Identification and estimation Our job seeker/firm pairwise design allows us to determine the extent to which our recommendation emails have led to an increase in the matching probability of a specific recommended job seeker/firm pair—which we label a targeting effect. Isolating this effect is important to determine whether recommender systems can be used by placement agencies to reduce matching frictions through tailored recommendations. Specifically, we define the *targeting effect* as the impact of a recommendation on the likelihood that a given match (i, j) occurs if firm j is recommended to job seeker i. We identify this targeting effect by comparing the match-level outcomes of treated job seekers in the firms recommended and non-recommended to them. In addition, we will refer to a *residual effect* if the match probability of treated job seekers' non-recommended firms increases compared to the same match probability of control job seekers. In other words the residual effect is defined as the increase in the likelihood that a match (i, j) occurs when job seeker i is in the treatment group rather than the control group, *in the absence of a recommendation for the pair* (i, j). The residual effect may combine the last three of the four mechanisms listed above: crowding out of non-recommended matches by recommended ones, increased search on the LBB platform, and overall increase in search effort.

Formally, the targeting and residual effects parameters are given by the following differences in expected outcomes:

$$\text{TARGETING}^{R=1} = E\left[Y_i^j(z_i = 1, r_i^j = 1) - Y_i^j(z_i = 1, r_i^j = 0) | R_i^j = 1\right] \tag{1}$$

$$\text{RESIDUAL}^{R=1} = E\left[Y_i^j(z_i = 1, r_i^j = 0) - Y_i^j(z_i = 0, r_i^j = 0) | R_i^j = 1\right] \tag{2}$$

where we enrich our counterfactual notations such that $Y_i^j(z_i, r_i^j)$ stands for the counterfactual employment outcome of job seeker/firm pair (i, j), which is a function of z_i that indicates if job seeker i belongs to the treatment group, and r_i^j that encodes recommended matches as opposed to non-recommended ones.

In the above formula both parameters are written as an average treatment effect on treated job seeker/firm matches (ATT), as they are conditional on the realized recommendation R_i^j. Because it is meant to maximize employment, our design is structured so as to give a higher recommendation probability $\pi_{i,j}$ to the job seeker/firm pairs that are more likely to give rise to a hiring, as these matches are both closer in terms of occupational distance and involve firms with relatively larger predicted hirings (Section 3.2 and Figure 1). As a consequence, pairs with high recommendation probabilities $\pi_{i,j}$ have higher baseline matching rates, as shown by Figure 2. However, our algorithm does not systematically target pairs with the most likely hirings; it takes into account possible congestion effects by recommending some firms in nearby occupations, trading off the lower congestion against potential occupational mobility costs. The targeting (ATT) parameter measures the extent to which the resulting recommendations improve the chances of treated job seekers at the selected firms.

It is interesting to compare these ATTs with the average treatment effects (ATE) that are the same parameters without the conditioning on $R_i^j = 1$. ATEs re-weight the matching success so as to be representative of the distribution of all possible matches and thus measures the effect of an intervention that would have made recommendations randomly (but among LLB firms only), without optimizing them.[20]

Both the targeting and residual effects are identified thanks to the randomization of recommendations, and estimated by inverse probability weighting (Horvitz-Thompson estimator). The weights account for the fact that recommendations are targeted on pairs with higher baseline matching likelihood: recommendations are orthogonal to potential outcomes after conditioning for (known) recommendation probabilities (see details in Appendix A.8).

Results Table VI, panel A, reports our estimates of the targeting and residual effects both in the ATT and ATE sense. The baseline probability that a given match (i, j) recommended by our design occurs is of course very low: about 0.04% (column (1)). As discussed previously, it is naturally higher than the baseline probability of occurrence of a random match, which is 0.01% (column (3)). We find a sizable and statistically significant targeting ATT effect in column (1): our recommendations increase the match probability by 18% ($\approx 0.00734/0.04716$). The residual effect is also positive, but lower and imprecisely estimated: there is no clear evidence that substitution effects compensate the targeting effects, neither can we strictly exclude that there are residual effects contributing to the overall, reduced form, effect in Table V. Overall, the targeting effect seems to be the main driver of increased job finding. The targeting effect in an ATE sense is much lower and not statistically significant, although it is rather large in comparison to the very low baseline match

[20] In practice, when computing parameters in the ATE sense, we restrict possible matches to job seekers and firms located in the same commuting zones, at occupational distance below 4.

FIGURE 2.—PAIRWISE BASELINE OUTCOME OF CONTROL JOB SEEKERS BY RECOMMENDATION PROBABILITY GROUP

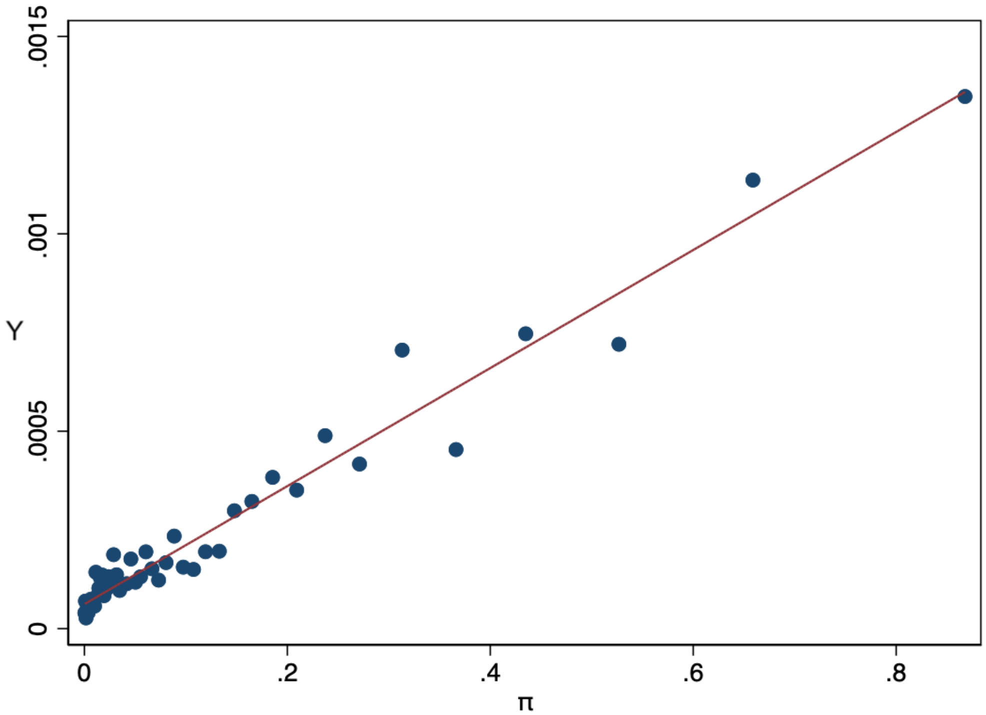

Notes: This figure displays a binned scatter plot of the baseline outcome of control job seekers' potential matches on the recommendation probability π. Formally, in the notations introduced in the main text and denoting by b_k the k-th bin of propensity score, each dot reports an estimate of $\mathrm{E}\left[Y_i^j \mid \pi_{ij} \in b_k, R_i^j = 0, Z_i = 0\right]$, which identifies (by design of our experiment, and assuming the bins are small enough to control effectively for relevant variations in the propensity score) $\mathrm{E}\left[Y_i^j(Z_i = 0, R_i^j) \,|\pi_{ij} \in b_k\right]$.

probability. It is not surprising that, even when recommended, less likely potential matches result in lower hiring rates.

This finding on the targeting effect is important because it formally implies that some job seekers sent additional applications to recommended firms, and that such applications were effective. However, they may remain prohibitively costly, even with the help of a recommender system such as the one implemented here, if applications have low success rates. The next section sheds light on this question, by decomposing the targeting effect into the rate at which job seekers send applications and the rate at which firms hire applicants.

TABLE VI

TARGETING AND RESIDUAL EFFECTS, AND APPLICATION EFFICIENCY

	Targeting (1)	Residual (2)	Targeting (3)	Residual (4)
	ATT		ATE	
A. Targeting and residual effects				
(a) Effect ($\times 100$)	0.00734	0.00403	0.00482	-0.00003
	(0.0033)	(0.00350)	(0.00477)	(0.00031)
	[0.03]	[0.25]	[0.31]	[0.91]
(b) Baseline ($\times 100$)	0.04176	0.03773	0.01203	0.01206
N	49,068,302	71,341,446	49,068,302	71,341,446
B. Application efficiency				
ρ: application rate	0.0728		0.0549	
μ: application efficiency	0.00696		0.00259	
$= \frac{1}{100} \cdot \frac{\text{(a)} + \text{(b)}}{\rho}$	(0.0008)		(0.0005)	
$\rho - \rho_0 = \frac{\text{(a)}}{\mu}$	0.0105		0.0186	

Notes: Panel A of this table presents estimates of the ATT and ATE for both the targeting and residual effects of the intervention at the dyad level (as defined in the main text). Panel B presents estimates of the application rate and application efficiency, also defined in the main text. Standard errors are clustered at the labor market level (CZ*occupations) and are reported in parentheses. Associated p-values are reported in square brackets. Coefficients and standard errors in lines (a) and (b) are reported in percentage points (estimates $\times 100$).

4.4. *Application efficiency: toward a measure of the potential of directing job search*

Identifying application efficiency Let us give a slightly more "structural" content to the counterfactual representation of the employment outcome: $Y_i^j(a_i^j, r_i^j)$, where $a_i^j = 1$ if job seeker i applies to firm j, and zero otherwise, and $r_i^j = 1$ if the pair (i, j) was recommended, as before. We now only have in mind treated job seekers ($Z_i^j = 1$), so we omit this dimension for simplicity. We also define the counterfactual application dummy variable $A_i^j(r_i^j)$. Also, we assume that only applicants can be hired, thus $Y_i^j(a_i^j = 0, r_i^j) = 0$. In that case:

$$E\left[Y_i^j|R_i^j=1\right]=E\left[Y_i^j|A_i^j=1,R_i^j=1\right]P\left[A_i^j=1|R_i^j=1\right]+0$$
$$=E\left[Y_i^j(a_i^j=1,r_i^j=1)|A_i^j(r_i^j=1)=1,R_i^j=1\right]P\left[A_i^j(r_i^j=1)|R_i^j=1\right].$$

For simplicity of exposition and table labeling, let us call $\mu=E[Y_i^j(a_i^j=1,r_i^j=1)|A_i^j(r_i^j=1)=1,R_i^j=1]$ and $\rho=P[A_i^j(r_i^j=1)|R_i^j=1]$. The former is the probability that a given application is selected by the firm for hiring, which we label *application efficiency*, estimated on the set of recommended dyads for which i applied to j. The latter is the probability to send an application to a recommended firm, on the set of recommended dyads (thus in an ATT sense). As we asked job seekers if they applied or not to each of the firm that we recommended to them, we observe ρ in the survey data, and we can compute μ as:

$$\mu=\frac{E\left[Y_i^j|R_i^j=1\right]}{\rho}.$$

This parameter allows us to assess the (average) efficiency of the applications sent by job seekers to the LBB firms suggested by the recommender system. As will appear in the tables, μ can be precisely estimated because the numerator is directly estimated as a simple average. Notice that we can also identify and estimate the average efficiency of applications sent by job seekers re-weighted by the distribution of all possible matches, rather than by the distribution of the recommended ones.[21] Formally, this corresponds to $E\left[Y_i^j(a_i^j=1,r_i^j=1)|A_i^j(r_i^j=1)=1\right]$, and it is the analog of μ in an ATE sense. Comparing estimates of these two parameters will allow us to document whether, and to what extent, our recommendation model encouraged a set of more effective applications.

Lastly, call $\rho_0=P[A_i^j(r_i^j=0)|R_i^j=1]$ the probability for i to send an application to j when the dyad has not been recommended (but computed on the distribution of recommended dyads, still in an ATT sense). We are also interested in $\rho-\rho_0$ as a measure of the effect of recommendations on sending applications. We do not observe ρ_0 however,

[21]In practice, we do so while still restricting our attention to firms hiring in occupations sufficiently connected to the occupation originally searched by job seekers (occupational distance below 4).

because we did not ask job seekers if they sent applications to firms that were not recommended to them (given the relatively small application rate, we would have needed to provide extremely long lists to measure this parameter adequately). But we can recover ρ_0 if we impose two restrictions. First, the exclusion restriction that $Y_i^j(a_i^j, r_i^j) = Y_i^j(a_i^j)$, $\forall(a,r)$: applicants to a job are no more likely to obtain the position if we have recommended the match than if we have not. Second, the homogeneity condition that $E[Y_i^j(a_i^j = 1, r_i^j = 1)|A_i^j(r_i^j = 1) = 1, R_i^j = 1] = E[Y_i^j(a_i^j = 1, r_i^j = 0)|A_i^j(r_i^j = 0) = 1, R_i^j = 1]$: the efficiency structure of applications if recommended is not different from that structure if not recommended (but computed on recommended dyads, $R_i^j = 1$).

In that case, with the new counterfactual notations, we can then re-write equation 1 as:

$$\begin{aligned}
\text{TARGETING}^{R=1} &= E\left[Y_i^j(a_i^j = 1, r_i^j = 1)|A_i^j(r_i^j = 1) = 1, R_i^j = 1\right] P\left[A_i^j(r_i^j = 1) = 1|R_i^j = 1\right] \\
&\quad - E\left[Y_i^j(a_i^j = 1, r_i^j = 0)|A_i^j(r_i^j = 0) = 1, R_i^j = 1\right] P\left[A_i^j(r_i^j = 0) = 1|R_i^j = 1\right] \\
&= \mu(\rho - \rho_0)
\end{aligned}$$

from which we can infer $\rho - \rho_0$.

Results Results are presented in Table VI, panel B, column 1. According to the survey, the average application rate to any recommended firm is 7.28%, which encompasses applications triggered by our recommendations, but also applications that would have occurred even without recommendations. Together with the probability that the match gets realized, this application rate implies that the application efficiency μ is 0.696%, very precisely estimated.[22] This is a relatively high success rate of applications, as it implies that one in every 143 applications is successful. Sending such applications is thus a reasonable strategy, as long as they are well targeted.

Nevertheless, in absolute value, the ATT targeting effect (panel A, column (1)) is relatively low: the average recommendation generates less than 0.01 percentage point more hirings. The reason is that the recommendations only have a very limited effect on the

[22]To facilitate the reading of the table, we present μ as a function of (a) and (b), but the actual estimation uses directly $E[Y_i^j|R_i^j = 1]$ and does not take the detour of using the estimated targeting effect.

application rate which increases by only one percentage point (a 14% increase). Our estimates imply that the baseline application rate ρ_0 is $6,2\%$, a rate that may seem high in the absence of recommendations but which can be explained by the fact that 23% of treated job seekers are using LBB (and even 20% in the control group), and that these firms are likely targets anyway. As the number of local LBB firms in a given occupation is limited (about 10 per occupation and 100 in neighboring occupations), it is not surprising that the baseline application rate of treated job seekers to recommended firms reaches 6%. What is more disappointing, however, is that recommendations do not boost application rates more strongly.

Outlining the potential of directing job search The implication of these findings is that the potential for improving the targeting of job seekers' applications is considerable, although a simple low-cost, email-based, intervention was not sufficient to increase applications by enough to generate substantial increases in hiring rates. Indeed, column (3) of Table VI computes application efficiency in an ATE sense, i.e. re-weighting the data so as to identify the average efficiency of applications sent by job seekers in the distribution of all possible matches: the resulting ATE application efficiency is 2.7 times ($\approx 0.00696/0.00259$) as low as the ATT μ, and it implies that on average only one out of 386 applications sent by a job seeker to firms taken from the population of all LBB firms would be successful. Ultimately, this difference between the application efficiencies measured in columns (1) and (3) suggests that any policy manipulating applications in a *directed* way—i.e. based on the matches that the model has selected to recommend—would be close to three times as effective at raising job finding rates as indiscriminate applications to LBB firms.

The difference in estimated application rates is also interesting: job seekers apply at a 7.28% rate for LBB firms recommended to them, but this drops to 5.49% when weighted by all LBB firms. It suggests that job seekers are more likely to apply to those firms that we recommend more often: on top of providing better hiring prospect, the firms identified by our system align relatively well with the preferences of job seekers.[23]

[23] Bied et al. (2023) underscore the fact that recommender systems may fail to maximize welfare if they do not take into account the preferences of job seekers in their design—preferences that may be revealed by the

For this potential to materialize, however, several conditions should be met. First, job seekers' applications should increase more strongly than in our experiment—this could be achieved by sending repeated emails, or involving case workers in communicating the recommendations. Secondly, the recommender system needs to identify a sufficient number of firms with a high untapped hiring potential. Thirdly, it needs to consider the possible switching costs and congestion effects generated by the redirection of job seekers to those firms. In the next two subsections, we turn to what our design allows us to learn on these two aspects.

4.5. *Heterogeneity of the targeting effect*

Our design made recommendations to maximize employment based on two key assumptions: that firms identified as hiring firms by the LBB algorithm would be more likely to hire when facing an increase in applications, and that recommending job seekers to broaden their search to tighter neighboring occupations could be beneficial despite occupational mobility costs. To shed light on these assumptions, Table VII provides the same decomposition as Table VI, splitting the sample by firm and local labor market characteristics. To avoid data mining, we restrict the analysis of heterogeneous treatment effects in Table VII to the three dimensions corresponding to the mechanisms embedded in our labor market matching model (see Section 3). As always, the differences across causal effects in the different columns need not be causal themselves, and may reflect other dimensions of heterogeneity across job seekers, firms or local markets. They are however suggestive that the mechanisms underlying our recommender system are at play.

Which firms is it more effective to recommend? In columns (3) and (4) of Table VII, we find that the targeting effect and the application efficiency are larger when recommending LBB firms with higher predicted hiring. Specifically, the application efficiency μ is about three times as large when the recommended firm has hiring forecasts above median compared to those below (0.00950 vs. 0.00277). As recommendations increase application rates for both types of firms by about one percentage point (0.01025 and 0.01172), targeting effects parallel application efficiencies: the targeting effect is about four times as large

application behavior of job seekers. In this case, we may find it reassuring that job seekers apply at relatively higher rates to the firms we recommend them.

when the firm has above median predicted hiring (0.01113 vs. 0.00284). Importantly, job seekers do not appear to be fully aware of this large heterogeneity in application efficiency: application rates in the absence of our recommendations (ρ_0) differ by less than 2 percentage points for firms with above vs. below median hiring forecasts. Overall, these results confirm the capacity of the LBB algorithm to identify firms with high hiring potential.[24]

Is broadening job search effective? Given that the LBB platform appears to identify firms with higher hiring potential, a naive use of its algorithm would be to systematically recommend the very top LBB firms to all job seekers. Such a strategy, however, would probably lead to overshooting, for two reasons. First, it would impose high mobility costs to job seekers searching in distant occupations. Second, it would create high levels of workers congestion at these top firms. Instead, our flexible matching model aims to recommend firms in local markets where tightness is high to more job seekers, while factoring in the costs of occupational mobility and congestion. Columns (5) and (6) suggest that targeting tighter markets is an efficient strategy.[25] The recommendations increase application rates by about one percentage point, irrespective of the tightness of the recommended local labor market. However, the estimates of μ implies that applications to less tight local markets are one-third less efficient than to tighter labor markets (p-value = 0.12). As a result, targeting effects are also about one-third higher when targeting a tighter market. However, for job seekers that are not searching in a tight occupation, this implies changing occupations: columns (1) and (2) consider the cost of that change, by comparing the effect of recommendations to the job seeker's own occupation with the effect of recommending a neighboring occupation. Recommendations increase application rates by about one percentage point, both in own and neighboring occupations. As application efficiency is one-third lower in a neighboring occupation (p-value of the difference = 0.12)—consistent with a loss in human capital taken into account by firms—targeting effects are about one-third lower when rec-

---

[24]Strictly speaking, columns (3) and (4) only make comparison between *bonnes boîtes* with higher and lower predicted hiring, rather than between *bonnes boîtes* and non-*bonnes boîtes* (firms that have too low predicted hiring to be advertised by LBB). However, these different categories are based on the same (continuous) variable, the predicted hiring from the LBB algorithm. Columns (3) and (4) suggest that these predictions capture relevant differences in hiring potential of firms in general.

[25]Tightness is measured as local labor markets' specific job finding rates. These job finding rates are recovered as the local labor markets fixed effects in a regression of job finding on control job seekers' characteristics.

ommending neighboring occupations. This suggests that, from the job seeker's perspective, at our equilibrium, the cost of changing occupations are not entirely offset by the benefit of applying to an occupation in a tighter market.[26]

To summarize, Table VII provides clear evidence on the potential of recommender systems to fruitfully redirect job search even *within* markets, as it highlights substantial heterogeneity in firms' application efficiency—especially with respect to their (predicted) hiring dynamics. Such heterogeneity is consistent with the existing literature on firm dynamics which has documented a wide variation in recruitment intensity, particularly as a function of firm employment growth (Davis et al., 2012, 2013). Yet it has received very little attention so far in studies of (directed) activation policies, where the emphasis has been on reorienting job search *across* occupations based on *market-level* hiring prospects (Belot et al., 2019). In our context, the effectiveness of such reorientation across markets is less clear. This may indicate that gaps in tightness between neighboring markets are not large enough, or that occupational mobility costs are too large.

[26]One needs to be cautious regarding the interpretation of this result. It is not inconsistent with previous evidence on the effectiveness of broadening job search (Belot et al., 2019, 2023, Altmann et al., 2023) for two reasons. Firstly, it sheds light on the importance of suggesting such occupational switching *only* when significantly better opportunities are available (e.g., going from a below median to an above median tightness labor market). Second, the application efficiency computed in column (1) need not be the counterfactual application efficiency that individuals who were recommended to apply at $d > 0$ would have faced. Indeed, such recommendations at $d > 0$ were made by our algorithm precisely when individuals were originally searching in relatively slack markets, which are demonstrated in column (3) to be markets where the average application efficiency is generally smaller.

TABLE VII

DECOMPOSITION OF TARGETING EFFECTS (ATT) AND APPLICATION EFFICIENCY BY TYPES RECOMMENDATIONS

	$d = 0$	$d > 0$	Pred. hiring below med.	Pred. hiring above med.	Mkt. tightness below med.	Mkt. tightness above med.
	(1)	(2)	(3)	(4)	(5)	(6)
A. Targeting and residual effects						
(a) ATT ($\times$100)	0.00802	0.00501	0.00284	0.01113	0.00648	0.00927
	(0.00287)	(0.00260)	(0.00198)	(0.00534)	(0.00282)	(0.00788)
	[0.08]	[0.05]	[0.15]	[0.04]	[0.02]	[0.24]
(b) Baseline ($\times$100)	0.0510	0.0162	0.0149	0.0617	0.0353	0.055
N	23,423,413	25,644,889	24,438,064	24,630,238	33,122,774	15,945,528
B. Application efficiency						
ρ: application rate	0.0817	0.0437	0.0631	0.0790	0.0718	0.0749
μ: application efficiency	0.00752	0.00499	0.00277	0.00950	0.00600	0.00890
$= \frac{1}{100} \cdot \frac{\text{(a)} + \text{(b)}}{\rho}$	(0.00095)	(0.00132)	(0.00056)	(0.00134)	(0.00085)	(0.00171)
[p-val. diff.]		[0.12]		[0.00]		[0.12]
$\rho - \rho_0 = \frac{1}{100} \cdot \frac{\text{(a)}}{\mu}$	0.01066	0.01004	0.01025	0.01172	0.01080	0.01042

Notes: Panel A presents estimates of the average targeting effect for different subgroups of potential job seeker/firm matches. Columns (1) and (2) split the sample of potential matches according to occupational distance ($d_{i,j} = 0$ and $d_{i,j} > 0$). Columns (3) and (4) split the sample of potential matches between matches involving firms with below or above median levels of predicted hirings by LBB. Columns (5) and (6) compare potential matches according to the tightness of the firm's local labor market (CZ $\times$ occupation). Panel B reports the corresponding application rates and application efficiencies (as defined in the main text). Standard errors are clustered at the labor market level (CZ $\times$ occupation) and are reported in parentheses. Associated p-values are reported in square brackets. Coefficients and standard errors in lines (a) and (b) are reported in percentage points (estimates $\times$100).

4.6. *Empirical evidence on congestion effects at the firm level*

As with any active labor market policy (ALMP), one cannot judge the potential of *directed*, recommender-based activation of job search without taking into account the possible congestion (or spillover) effects it may generate. It has been widely documented that such equilibrium effects can mitigate the direct effect of ALMPs (Crépon et al., 2013, Lalive et al., 2012). In our setting, the risk of a large increase in applications to a subset of firms, even if well chosen, is to generate (firm-level) congestion, whereby the probability that a firm chooses any applicant (μ) may decrease with the number of applicants. This will happen if firms do not endogenously create enough positions when the labor supply they face increases. Job creation by firms could in fact react (i) in full proportion to the number of applications received ("no congestion" case) or (ii) less than proportionally (as in our matching model generating our recommendations). In the latter case, we should observe that the likelihood of accepting any applicant decreases as we send more applications to firms.

Evidence on firm-level congestion Our experimental design allows to explore empirically the presence of congestion effects. Unlike most results on labor market spillover effects that estimate externalities at the market level, we can measure congestion at the more granular level of the firm.[27] Indeed, firms were randomly assigned to the "few" and

[27] We explored two alternative approaches to the estimation of externalities and found that the resulting estimators are too imprecise to be informative in our context. First, we followed Crépon et al. (2013) using as "super controls" the commuting zones that had (randomly) been excluded from our experiment. However, given the small direct effects of the intervention on treated job seekers, it is not surprising that a comparison at the aggregate level across commuting zones is not able to detect displacement effects. Second, we used the variation in the exposure of control job seekers to the treatment that is induced by the Bernoulli trial, following Hu et al. (2022). However, the resulting variation is limited in our context, and, in the absence of further assumptions on the interference structure, the estimates of the " average indirect effect" that we find are very imprecise. This lack of precision in leading approaches in the literature prompts us to focus on within-firm congestion, for which our design intentionally generates variation.

"many" recommendations treatment arms, as explained at the end of Section 3.2. Firms in the "many" recommendations treatment arm get recommended twice as much as firms allocated to the "few" recommendations group: firms in the "few" treatment arm were recommended to 40.7 job seekers on average; firms in the "many" treatment arm are recommended to 85.3 job seekers. This provides a random variation to assess the congestion effect of increasing the number of (job seekers) applications directed to a given firm by about 5%.[28] The result of this exercise is presented in Table VIII where we compute the probability of hiring an applicant μ separately in firms belonging to each treatment arm. This probability does decrease when the number of recommendations increases, the difference between the two groups being significant at the 10% level.

Vindicating our matching model of the labor market Whereas the nature of our experiment (and in particular the lack of statistical power) does not allow us to pin down a precise value for the congestion elasticity, we see this result as an important tale of caution when designing automated recommender systems. The potential gains stemming from large variations in local application efficiencies should be balanced against potentially large congestion effects. Our own design offers a way of resolving this difficult trade-off on the basis of economic reasoning, through parsimonious modeling of the labor market—and in particular of the firm hiring process. As a reminder, the key elements of this flexible model of the labor market are (i) heterogeneous propensities to hire (conditional on the number of applications per vacancy), (ii) a penalty for occupational switching in the probability of being hired (accounting for human capital losses), and (iii) partial congestion in the hiring

[28]There are (on average) 32 job seekers in each market, and 128 job seekers in neighboring markets (with occupational distance below 4). Since these 160 job seekers apply at a rate of 6.3% in the absence of any intervention (cf. Table VI), this represents 10 applications per firm. Through our experiments, we boosted by 1 percentage point the applications of 40.7 job seekers toward firms in the "Few" arm ($+0.4$ application), and 85.3 job seekers toward firms in the "Many" arm ($+0.85$ application). Hence an increase in the number of applications per firm between the "Few" and "Many" arms of $(0.85-0.4)/(10+0.4) \approx 4.3\%$.

process (modeled as a less-than-proportional response of the firm's hiring to the increase in the number of applications). Our results show that each of these elements is relevant from an empirical point of view, suggesting that our matching provides a useful framework to design effective recommender-based activation policies.

TABLE VIII

THE EFFECT OF THE NUMBER OF RECOMMENDATIONS ON THE EFFICIENCY OF APPLICATIONS (μ)

	Few rec. (1)	Many rec. (2)
μ: application efficiency	0.00765	0.00663
	(0.00097)	(0.00079)
[p-val. diff.]		[0.09]

Notes: This table presents estimates of the application efficiency (as defined in the main text) in firms that randomly received "Few" vs. "Many" recommendations. Standard errors are clustered at the labor market level (CZ $\times$ occupations) and are reported in parentheses. The p-value for the test of no difference between the two estimates is reported in square brackets in column (2).

5. CONCLUSION

Building upon an existing job search platform operated by the French PES, we show that recommender systems have the potential to improve job seekers' labor market outcomes by redirecting job search effort toward hiring firms and tighter occupations. We generate specific recommendations using a flexible model of the labor market that seeks to optimize the potential employment rate. Our empirical results demonstrate that such a model is useful to take full advantage of the heterogeneity in firms hiring behavior, while factoring in mobility costs and congestion.

But our study uses an encouragement design: e-mailing treated job seekers with firm-specific job search advice. Such designs typically have limited take-up, and our study is no exception. In that context, the large scale of the experiment is key for two reasons. First, it allows us to detect small effects with sufficient precision. Second, it shows that a realistic, low-cost intervention, can have real-life effects. It remains however the case that effects are small, when expressed in terms of job finding rates. While this does not prevent the policy to be very likely cost effective (given its very low cost), it begs the question of whether features of the intervention could be enhanced to increase impact. Recent work by Altmann et al. (2023) demonstrate the effectiveness of integrating simpler occupation-based redirections in widely used platforms. They also document large displacement effects. Therefore, getting a large share of job seekers to use a congestion-aware recommender system such as ours—exploiting the full potential of job search redirections documented in this paper at the firm level, both across *and* within occupations— remains a promising and important avenue for future research.

BIBLIOGRAPHY

ALTMANN, STEFFEN, ANITA GLENNY, ROBERT MAHLSTEDT, AND ALEXANDER SEBALD (2023): "The Direct and Indirect Effects of Online Job Search Advice," Working Paper. [7, 32, 37]

ANDERSON, MICHAEL L. (2008): "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects," Journal of the American Statistical Association, 103 (484), 1481–1495. [61]

ATHEY, SUSAN, JULIE TIBSHIRANI, AND STEFAN WAGER (2018): "Generalized random forests," Annals of Statistics (forthcoming). [57]

AUTOR, DAVID H (2001): "Wiring the Labor Market," Journal of Economic Perspectives, 15 (1), 25–40. [7]

BELOT, MICHELE, DIDIER FOUARGE, PHILIPP KIRCHER, PAUL MULLER, AND SANDRA PHLIPPEN (2023): "Stimulating Occupational Mobility among Unemployed Job Seekers," Working Paper. [7, 32]

BELOT, MICHELE, PHILIPP KIRCHER, AND PAUL MULLER (2019): "Providing Advice to Job Seekers at Low Cost : An Experimental Study on On-Line Advice," Review of Economic Studies. [3, 7, 32]

——— (2022a): "Do the Long-term Unemployed Benefit from Automated Occupational Advice during Online Job Search?" Working Paper. [7]

——— (2022b): "How wage announcements affect job search – a field experiment," American Economic Journal - Macroeconomics. [3]

BEN DHIA, AÏCHA, BRUNO CRÉPON, ESTHER MBIH, BERTILLE PICARD, AND VINCENT PONS (2022): "Can a Website Bring Unemployment Down? Experimental Evidence from France," Working Paper. [7]

BIED, GUILLAUME, PHILIPPE CAILLOU, BRUNO CREPON, CHRISTOPHE GAILLAC, ELIA PERENNES, AND MICHÈLE SEBAG (2023): "Designing labor market recommender systems: the importance of job seeker preferences and competition," Working Paper. [7, 8, 29]

CHERNOZHUKOV, VICTOR, MERT DEMIRER, ESTHER DUFLO, AND IVAN FERNANDEZ-VAL (2018): "Generic machine learning inference on heterogenous treatment effects in randomized experiments," . [21, 59, 60]

CHERUBINI, MAURO, MATTHIAS KLIEGEL, RAFAEL LALIVE, MICHELE PELLIZZARI, AND DORIANA TINELLO (2023): "Improving Job Search: Matching and Cognition," . [7]

CRÉPON, BRUNO, ESTHER DUFLO, MARC GURGAND, ROLAND RATHELOT, AND PHILIPPE ZAMORA (2013): "Do labor market policies have displacement effects? Evidence from a clustered randomized experiment," Quarterly Journal of Economics, 531–580. [34]

DAVIS, STEVEN J., JASON R. FABERMAN, AND JOHN C. HALTIWANGER (2012): “Recruiting Intensity during and after the Great Recession: National and Industry Evidence,” American Economic Review: Papers & Proceedings. [3, 6, 8, 32]

——— (2013): “The Establishment-Level Behavior of Vacancies and Hiring,” Quarterly Journal of Economics. [3, 8, 32]

HENSVIK, LENA, THOMAS LE BARBANCHON, AND ROLAND RATHELOT (2023): “How can AI improve search and matching? Evidence from 59 millions personalized job recommendations,” Working Paper. [7]

HORTON, JOHN J (2017): “The Effects of Algorithmic Labor Market Recommendations : Evidence from a Field Experiment,” Journal of Labor Economics, 35 (2). [3, 7]

HU, YUCHEN, SHUANGNING LI, AND STEFAN WAGER (2022): “Average direct and indirect causal effects under interference,” Biometrika, asac008. [34]

KIRCHER, PHILIPP (2020): “Search Design and Online Job Search–New Avenues for Applied and Experimental Research,” Labour Economics, 101820. [3]

——— (2022): “Job Search in the 21st Century,” Journal of the European Economic Association, 20 (6), 2317–2352. [3, 4, 6]

KUHN, PETER J (2014): “The internet as a labor market matchmaker: how effective are online methods of worker recruitment and job?” IZA World of Labor, (May), 1–10. [3, 7]

KUHN, PETER J AND HANI MANSOUR (2013): “Is internet job search still ineffective?” The Economic Journal, 124, 1213–1233. [3, 7]

KUHN, PETER J AND MIKAL SKUTERUD (2004): “Internet Job Search and Unemployment Durations,” American Economic Review, 94 (1), 212–232. [3, 7]

LALIVE, RAFAEL, CAMILLE LANDAIS, AND JOSEF ZWEIMÜLLER (2012): “Market Externalities of Large Unemployment Insurance Extension Programs,” American Economic Review. [34]

LE BARBANCHON, THOMAS, MADDALENA RONCHI, AND JULIEN SAUVAGNAT (2023): “Hiring Difficulties and Firm Growth,” Working Paper. [8]

MARINESCU, IOANA AND ROLAND RATHELOT (2018): “Mismatch Unemployment and the Geography of Job Search,” American Economic Journal: Macroeconomics, 10 (3), 42–70. [14]

ROSENBAUM, PAUL R. AND DONALD B. RUBIN (1983): “The central role of the propensity score in observational studies for causal effects,” Biometrika, 70 (1), 41–55. [17]

VAN DER KLAAUW, BAS AND HEIKE VETHAAK (2022): "Empirical Evaluation of Broader Job Search Requirements for Unemployed Workers," IZA Working Paper 15698. [8]

YADLOWSKY, S., S. FLEMING, N. SHAH, E. BRUNSKILL, AND S. WAGER (2021): "Evaluating Treatment Prioritization Rules via Rank-Weighted Average Treatment Effects," Working Paper. [21, 57, 59]

ONLINE APPENDIX

A.1. *Context*

FIGURE A1.—LBB'S HOME PAGE

FIGURE A2.—LBB'S RESEARCH RESULTS PAGE

FIGURE A3.—LBB'S FIRM CONTACT INFORMATION PAGE

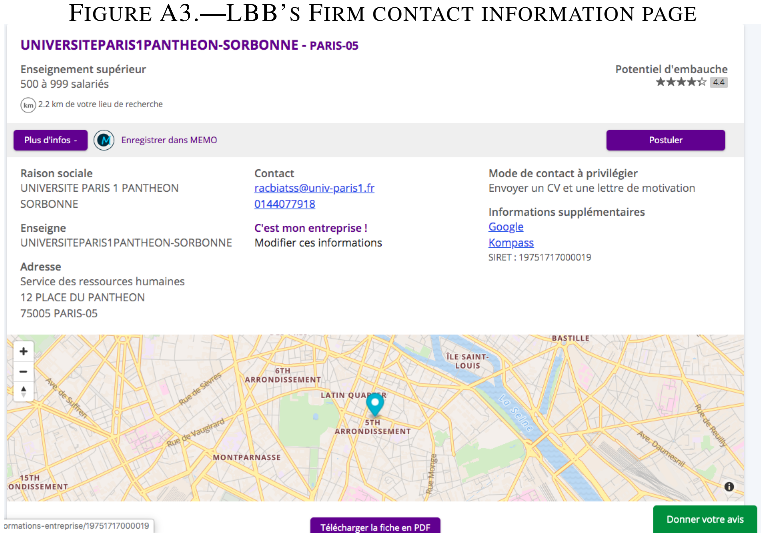

FIGURE A4.—EMAIL SENT TO TREATED JOB SEEKERS

Bonjour M. Zuber,

Vous êtes inscrit à Pôle emploi et avez déclaré rechercher un emploi dans la catégorie : « Sommellerie ».

Savez-vous que 7 entreprises sur 10 examinent des candidatures spontanées avant de se décider à publier une offre d'emploi ?

La Bonne Boite, un service de Pôle emploi, a repéré des entreprises que votre profil pourrait intéresser.

En voici une susceptible de rechercher un profil proche du vôtre :

- GSF MERCURE

Vous pouvez leur envoyer une candidature spontanée.

En cliquant sur ce lien, vous pourrez contacter l'entreprise grâce aux coordonnées qui s'affichent ou en utilisant l'outil de candidature en ligne « **postuler** » lorsque celui-ci est disponible.

Vous avez également la possibilité de retrouver d'autres entreprises sur le site La Bonne Boite

En vous souhaitant une pleine réussite dans votre recherche d'emploi.

A.2. *Balancing tests*

TABLE A1
BALANCE TABLE FOR JOB SEEKERS IN TREATED CZ.

	(1)		(2)		(3)	
	Control		Treated		Difference	
Gender	0.450	(0.498)	0.451	(0.498)	0.001	(0.001)
Age	38.944	(12.052)	38.975	(12.043)	0.030	(0.029)
Diploma	0.608	(0.488)	0.608	(0.488)	-0.000	(0.001)
Experience (y)	6.917	(8.198)	6.920	(8.202)	0.003	(0.019)
Unemployment spell (m)	21.258	(24.724)	21.313	(24.807)	0.055	(0.059)
Predicted exit rate	0.207	(0.072)	0.207	(0.072)	0.000	(0.000)
Predicted tightness	0.392	(0.660)	0.391	(0.666)	-0.000	(0.002)
Observations	266,740		533,557		800,297	

Note: Standard errors are displayed in parentheses.

A.3. *Occupational distance and observed transitions*

Both the PES and LBB use the same 532-occupations ROME classification (“Répertoire Opérationnel des Métiers”). used by the PES when asking job seekers their desired occupation, and by LBB to compute hiring predictions. In addition, we take advantage of PES’ expert knowledge on possible transitions to build a simple measure of occupational distance. More precisely, for every single occupation, the PES lists a set of neighbor occupations which are deemed close enough in terms of required skills for job seekers to transition to without any further training. We use these neighboring occupations to build an occupational graph where each occupation is connected to its listed neighboring occupations. As the closeness of occupations is not necessarily symmetric (occupation A neighboring occupation B does not entail that occupation B neighbors occupation A), the underlying occupational graph is a directional one. Finally we use this occupational graph to measure the relative closeness of any two occupations. To do this we compute the shortest path linking any two occupations and take this shortest path as our main measure of occupational distance. With this methodology 6.20% of occupations end up isolated, the average occupational distance between any two connected occupations, measured by the number of intermediary nods, is 7.11 and occupations are on average connected to 3.34 immediate neighbor occupations. As shown in Figure A5 of Appendix A.3, our measure of occupational distance correlates well with occupational transitions observed in the French data over the 2008/2012 period. Importantly, by limiting ourselves to PES' original definition of "close" occupations we only would have covered 15% of observed transitions. By extending our measure of occupational distance to pairs which were not previously ranked we are able to cover 83% of observed occupational transitions, hence giving a much more comprehensive view of the underlying occupational structure of the French labor market.

FIGURE A5.—MEAN OCCUPATION DISTANCE VS OBSERVED RANK IN OCCUPATIONAL TRANSITIONS

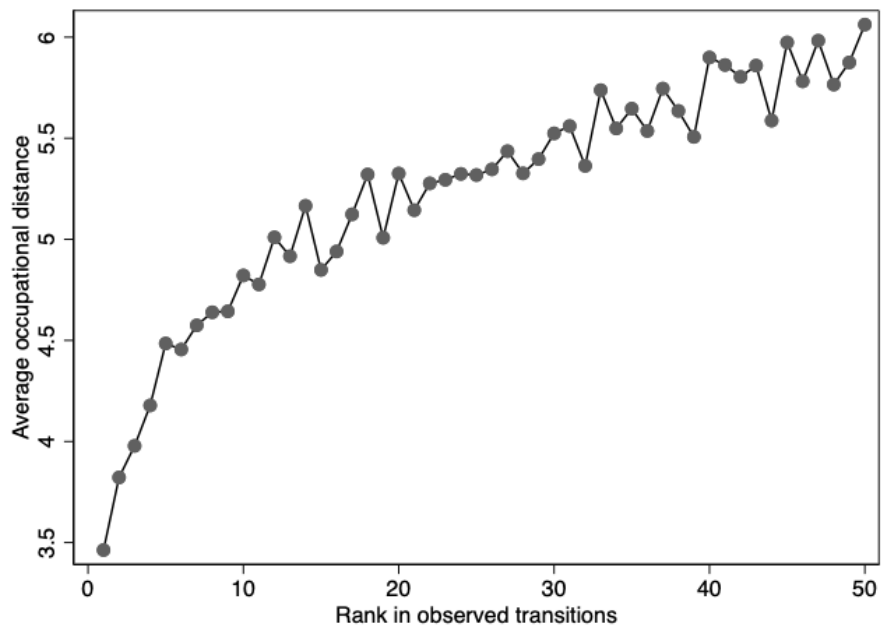

Note: This graph is constructed by ranking occupational transitions according to their frequency within each origin occupation and then computing the mean occupational distance of these transition in each rank category. In other words, across all origin occupations, destination occupation ranked first in terms of transitions were located at an average occupational distance of 3.5. Data on occupational transitions are constructed from the FHDADS panel covering the 2008-2012 period. We are constrained to this rather short period because prior to 2008 the DADS did not record a 4-digit occupation. An occupational transition from A to B is defined as a job seeker looking for a job in occupation A finding a job in occupation B. While the search occupation A is coded in the ROME classification, the destination occupation B is coded according to the PCS classification used in DADS files. We translate the PCS classification into the ROME one by using the ROME-FAP-PCS matching provided by the French unemployment agency as well as each ROME's distribution of educational attainments among job seekers observed in our pre-treatment data. In total this graph is constructed from 1,092,233 individual transitions over the 2008-2012 period

A.4. *Commuting zones and local labor markets*

A.4.1. *Commuting Zones*

For administrative purposes the PES divides the french territory into 404 commuting zones ("bassins d'emploi"). A commuting zone is a geographical space where most of the population lives and works. In other words, most people do not leave this area to go to their place of work. Both job seekers and firms are thus mapped to an specific commuting zone through their zip code. These areas have an average population of $160,000$ and are spread over an average radius of 20.3km.[29] Finally, and consistent with France's unemployment rate, there are on average $13,467$ job seekers in each commuting zone.

For this experiment 94 commuting zones out of the 404 initial ones were selected. We leave the 310 remaining commuting zones untouched for a future experiment guided by the learnings of this one. Nevertheless this experiment remains a large-scale experiment with more than 1.2 million job seekers and 750 thousand firms involved. The 94 commuting zones of our interest are randomly selected from the pool of commuting zones. Table A2 shows the main characteristics of commuting zones selected for the experiment (column 1) and commuting zones not selected for the experiment (column 2). We observe that characteristics between those groups are balanced and therefore our sample is representative of the entire France.

[29]We miss data for one commuting zone which regroups Saint-Martin and Saint-Barthélémy.

TABLE A2

COMMUTING ZONES' STATISTICS

	(1)	(2)	(3)
Variable	Selected Zone	Non Selected Zone	(2)-(1)
Surface (m2)	182507.453	150871.219	-31636.240
	(423423.031)	(200091.297)	(31,679.127)
Population	154650.000	161688.672	7,038.673
	(133044.750)	(196349.313)	(21,628.875)
Number of Unemployed	12,870.830	13,648.951	778.122
	(12,109.896)	(17,855.393)	(1,966.694)
Unemployment Ratio	0.079	0.081	0.002
	(0.017)	(0.019)	(0.002)
Number of Hiring Firms	7,985.681	8,512.371	526.690
	(9,362.619)	(15,645.074)	(1,699.878)
Tightness	0.623	0.585	-0.038
	(0.402)	(0.241)	(0.034)
Observations	94	310	404

Standard errors in parenthesis.

A.4.2. *Local Labor Markets*

Upon registrating with public employment services, job seekers are asked to fill in a certain number of personal information including their desired occupation. As one's desired occupation is not, however, a required information we drop job seekers whose search occupation appears as missing in our data. Job seekers who choose to register a desired occupation can select one occupation from the 532 options given in the "ROME" classification of occupations used by french unemployment services[30]). We define a local labor market

[30]ROME stands for "Répertoire opérationnel des métiers": Operational directory of occupations.

as the intersection between commuting zones and occupations. In France there are 404 CZ ands 532 occupations, which makes $404 \times 532 = 214928$ local labor markets. Among these potential labor market only 174733 turn up with a least one job seeker or one active establishment. On average a local labor market is populated by 31 job seekers and 19 establishments which total 12 predicted hirings. The mean predicted hirings to job seekers ratio is 0.31. This ratio can be thought of as the predicted tightness of our local labor markets.

A.5. *Correlating predicted and realized hirings*

FIGURE A6.—REALIZED HIRINGS AMONG UNEMPLOYED JOB SEEKERS OVER THE 30/09/2019-13/03/2020 PERIOD VS LBB'S PREDICTED HIRINGS AS OF 11/08/2019 (IN LOGS)

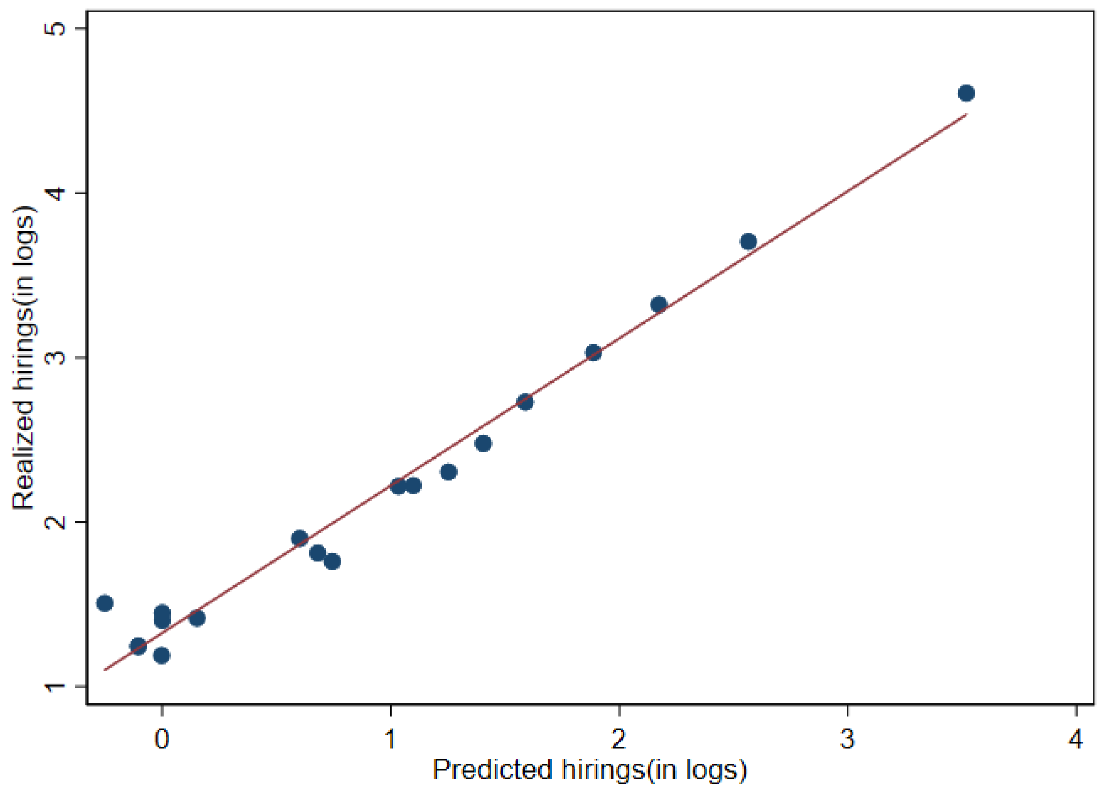

Note: Correlation of the number of predicted hirings per establishment and the number of realized hirings. LOG(REALIZED HIRINGS) $= 1.33(0.0053) + 0.89(0.0039) \times$ LOG(PREDICTED HIRINGS), $R^2 = 0.37$

A.6. *A flexible model of worker/firm matches*

A commuting zone is populated by I workers (indexed i) and J firms (indexed j). We denote firm j's predicted hirings or "vacancies" as V^j. For simplicity, note the distance

$d_i^j = d_{h(i)h(j)}$ as the occupational distance between job seeker i's own occupation $h(i)$ and the occupation $h(j)$ for which firm j has positive predicted hirings.[31]

Recommendation probability. At the outset of the experiment we start by fixing the total number of recommendations which will be received by each job seeker.[32] Denoting this number by N_i, the total number of recommendations we need to generate is:

$$N = \sum_{i \leq I} N_i.$$

In practice we repeatedly draw these N_i recommendations from a worker specific generalized Bernoulli distribution over all possible firms with positive predicted hirings. Our statistical model of worker/firm matches should be rich enough to solve for the set of optimal generalized Bernoulli non-negative probability weights

$$0 \leq p_i^j \leq 1$$

verifying

$$\sum_{j \leq J} p_i^j = 1$$

where p_i^j is the probability to recommend firm j to worker i in each single draw of the generalized Bernoulli distribution. Taking N_i as given, the probability to recommend firm

[31] In the case where a firm is predicted to hire in several occupations we take d_i^j to be the minimum distance between job seeker i's search occupation and firm j's hiring occupation.

[32] We randomly send up to eight recommendations, see Section 2.2.

j to job seeker i at least once is given by:

$$P(R_i^j = 1) = 1 - (1 - p_i^j)^{N_i}$$

where the random variable R_i^j takes the value 1 if we recommend firm j at least once and 0 otherwise.

Expected number of matches. Let Y_i^j denote the random variable which takes the value 1 if job seeker i is eventually hired by firm j. Our objective is to select the distribution of worker specific recommendations so as to maximize the expected number of matches in the economy :

$$Y = \mathbb{E}\left[\sum_{i,j} Y_i^j\right]$$

which can be rewritten as:

$$Y = \sum_{i,j} \mathbb{E}[Y_i^j | R_i^j = 1] \times [1 - (1 - p_i^j)^{N_i}] + \mathbb{E}[Y_i^j | R_i^j = 0] \times [1 - p_i^j]^{N_i}.$$

In order to concentrate on the effect of targeted recommendations we normalize all default outcomes $\mathbb{E}[Y_i^j | R_i^j = 0]$ to zero. Under this normalization our main object of interest is worker i's probability of being hired in firm j conditional on being recommended to apply to this position.

Job seeker's application strategy. On the worker side, we assume that each job seeker i may look for a job in his origin occupation as well as in neighboring occupations. Each worker is characterized by an idiosyncratic distaste for occupational distance $\rho_i \in (0, 1)$. Conditional on receiving a recommendation to apply to firm j, $R_i^j = 1$, we assume that worker i applies to firm j with probability

$$P(A_i^j = 1 | R_i^j = 1) = \rho_i^{d_i^j}$$

where the random variable A_i^j takes the value 1 if worker i applies to firm j and 0 otherwise.

Firm's hiring strategiy. Given workers' application behavior, firm j will on average receive

$$A^j = \sum_i \rho_i^{d_i^j} R_i^j$$

applications.[33] Upon receiving these applications, firm j randomly selects a proportion $q^j \in (0,1)$ of them. We assume that this proportion of successful application continuously depends on the ratio of received application A^j to predicted hirings V^j that we observe empirically. Let $\theta^j = A^j/V^j$ denote this measure of firm-level slackness, we set the screening rate q^j to:

$$q^j = q_j(\theta^j)$$

where q_j is a firm specific screening function verifying $q_j \in (0,1)$, $q_j' \leq 0$, $q_j(0) = 1$, and $q_j(+\infty) = 0$.[34] This firm-specific screening friction is key to the model: If it were not present, there would be no congestion at the firm level, and no social cost of directing more applicants to a given firm.

[33] Consistent with the normalization above, workers' probability to apply to an un-recommended firm is 0.

[34] See Appendix A.6.1 for further details on the shape of this function.

Conditional on applying to j, a worker can expect to be interviewed with probability: [35]

$$\tilde{q}^j = \mathbb{E}[q_j(\theta^j)].$$

Finally, screened applicants go through a final step in which the firm decides to hire or reject each applicant based on occupational distance. We denote each firm's distaste for occupational distance ρ_j and, as in the worker case, assume that each screened applicant's probability of being hired is given by $\rho_j^{d_i^j}$.

Both that distaste and a parameter of function $q_j(.)$ that determines a firm's efficiency at screening (see Appendix Appendix A.6.1) are randomly allocated to firms. It generates random exposure of firms to the quantity and quality (occupational distance) of applicants.

Summing up the hiring process. We can break down our model into the following steps:

1. We recommend firm j to worker i.
2. Worker i who is more or less averse to occupational distance d_i^j applies to firm j with probability $\rho_i^{d_i^j}$.

[35] Because in general the function $q_j(.)$ is non linear, we approximate this expectation through a second-order Taylor expansion:

$$\tilde{q}^j \sim q_j(\mathbb{E}[\theta^j]) + \frac{\mathbb{V}[\theta^j]}{2}\frac{\partial^2 q_j}{\partial \theta^2}(\mathbb{E}[\theta^j])$$

where $\mathbb{E}[\theta^j]$ and $\mathbb{V}[\theta^j]$ can be computed explicitly. Indeed, under the assumption we made on workers' application process we know that:

$$\mathbb{E}[\theta^j] = \frac{\mathbb{E}[A^j]}{V^j} = \frac{\sum_i \rho_i^{d_i^j}(1-(1-p_i^j)^{N_i})}{V^j}$$

and

$$\mathbb{V}[\theta^j] = \frac{\mathbb{V}[A^j]}{(V^j)^2} = \frac{\sum_i \rho_i^{d_i^j}[1-(1-p_i^j)^{N_i}][1-\rho_i^{d_i^j}[1-(1-p_i^j)^{N_i}]]}{(V^j)^2}$$

3. Firm j skims through the applications it receives and randomly decides to look more deeply into q^j of them.
4. Firms review selected applications and decide whether or not to hire each reviewed applicant according to occupational distance. Each screened applicant is hired with probability $\rho_j^{d_i^j}$.

Probability of hiring. The probability that worker i is hired by firm j in our experimental setting is:

$$\mathbb{E}[Y_i^j | R_i^j = 1] \sim \rho_j^{d_i^j} \times \tilde{q}^j \times \rho_i^{d_i^j}.$$

Substituting for $\mathbb{E}[Y_i^j | R_i^j = 1]$ in the expression for Y (the expected total number of matches created by our intervention) gives:

$$Y \sim \sum_{i,j} \rho_j^{d_i^j} \times \tilde{q}^j \times \rho_i^{d_i^j} \times [1 - (1 - p_i^j)^{N_i}]$$

which is a non-linear function of the Bernoulli weights p_i^j that are central to our experimental design.[36]

The optimal recommendations. The problem of the central planner is to maximize Y over the space of possible worker specific distribution of recommendations. This problem has dimensionality of about 2 millions per commuting zone, which is of course too large to solve by brute force. To reduce the dimensionality of the problem we parameterize p_i^j using available information on workers and firms. Denote X_i^j the vector of worker/firm

[36]Notice that $\tilde{q}^j$ is an implicit function of the other parameters in this expression, including $p_{i,j}$, and the V^j's that are given by the data.

characteristics that will be used to generate p_i^j. We assume that:

$$p_i^j = \frac{\exp(\beta X_i^j)}{\sum_j \exp(\beta X_i^j)}$$

Hence the dimensionality of the problem is reduced to the number of worker/firm characteristics so that the maximization problem boils down to:

$$\max_\beta \sum_{i,j} \rho_j^{d_i^j} \times \tilde{q}^j \times \rho_i^{d_i^j} \times [1 - (1 - \frac{\exp(\beta X_i^j)}{\sum_j \exp(\beta X_i^j)})^{N_i}].$$

The vector X includes: worker and firm level observed characteristics (the firm predicted hirings V^j, and the worker/firm occupational distance d_i^j) and the parameters that have been randomly allocated to workers and firms (ρ_j, ρ_i, and the shape parameter of the screening function $q_j(.)$, called m_j, see Appendix A.6.1). The optimal parameter β will in particular be sensitive to the occupational distribution of job seekers and firms within each commuting zone. In the case where job seekers and firms would operate in very different occupations, large aggregate gains should be expected from reallocating workers across occupations, so that the optimal β would put little negative weight on occupational distance in forming pairwise worker/firm recommendations. The exact opposite occurs if workers and firms are evenly distributed across the occupational space.

Finally, the probability that we recommend firm j to worker i is

$$\pi_{i,j} \equiv P(R_i^j = 1) = 1 - (1 - (p_i^j)^*)^{N_i}$$

where $(p_i^j)^*$ denotes the optimal p_i^j.

A.6.1. *Choice of the screening technology:*

More specifically we choose to parametrize our screening function q^j as:

$$q_j(\theta^j) = \frac{1}{[1 + (\frac{\theta^j}{\Gamma m_j \bar{\theta}_j})^{\gamma}]^{1/\gamma}}$$

where $\gamma > 1$ and Γ are constants verifying:

$$\Gamma = (\frac{\gamma - 1}{2})^{-1/\gamma}$$

where m_j is a firm specific constant which interpret as screening effeiciency parameter, and where $\bar{\theta}_j$ denotes the local slackness ratio in firm j's hiring occupations. This local slackness ratio is defined as the ratio of possible recommendations present in the neighborhood of firm j to the total number of hirings predicted in firm's j hiring occupations. Formally:

$$\bar{\theta}_j = \frac{\sum_i \rho_i^{d_{i,j}} N_i}{\sum_h V^{j,h}}$$

For $\gamma > 1$ this function is monotonous in $\theta^j = A^j / V^j > 0$ and verifies:

$$q^j(0) = 1$$

$$q^j(+\infty) = 0$$

What's more q^j has an inflection point at $m_j \theta_j$ so that according to the value of m_j, firm's j congestion effect will start to quick in either before ($m_j = m_j^L < 1$) or after ($m_j = m_j^H >$

1) the number of recommendations sent to j relative to its predicted hirings (i.e A^j/V^j) reaches the local slackness ratio θ_j. In practice we select $m^L = 0.5$, $m^H = 1.5$ and $\gamma = 3$.

A.6.2. *Implementation of the recommendation design*

We randomly assign to each treated job seeker and firm, a value of the key parameters of the model of Section A.6: $\rho_i \in \{0.82, 0.94\}$ the worker's distaste for occupational distance; $\rho_j \in \{0.82, 0.94\}$, the firm's distaste for occupational distance; and a parameter that determines the efficiency of the firm's screening function m_j.[37] Using also the observed V^j and d_i^j, we solve numerically for the optimal values of β, separately in each of the 94 commuting zones. We then compute the probabilities $\pi_{i,j} = P(R_i^j = 1|X_i^j) = P(R_i^j = 1|\rho_i, \rho_j, m_j, V^j, d_i^j)$, and proceed to draw as many job seeker/firm recommendations as needed, using the generalized Bernoulli distribution described in Section A.6.

A.7. *Heterogeneity analyses: additional results*

Agnostic, machine learning based heterogeneity analyses Here, we report the results of two separate methodologies aiming at detecting treatment effect heterogeneity without any prior on which observables might be predictive of such heterogeneity.

The first approach reported is the one developed by Yadlowsky et al. (2021). Their methodology consists in fitting on a random half of the data a model of conditional average treatment effects

$$\tau(X_i) \equiv \mathrm{E}\left[Y_i(1) - Y_i(0)|X_i\right]$$

using (for instance) causal forests (Athey et al., 2018). Then, on the second random half of the data, one can explore the extent to which the intervention would have achieved a

[37]All those parameters are determined by a stratified randomization with probability 0.5, that uses the same strata as for the treatment status of job seekers and firms respectively, see section 3.1.

larger average treatment effect if it had been targeted a fraction q of job seekers with the largest predicted CATEs $\hat{\tau}(X_i)$. This fraction q is reported in the x-axis of Figure A7, while the y-axis reports (as defined above) the Targeting Operator Characteristic (TOC). Here, it corresponds to the gap between the ATE among this fraction q of targeted job seekers with the largest predicted CATEs and the overall ATE from treating everyone:

$$\mathrm{TOC}(q) \equiv \mathrm{E}\left[Y_i^{\mathrm{ST}}(1) - Y_i^{\mathrm{ST}}(0) | \hat{\tau}(X_i) > F^{-1}_{\hat{\tau}(X_i)}(1-q)\right] - \underbrace{\mathrm{E}\left[Y_i^{\mathrm{ST}}(1) - Y_i^{\mathrm{ST}}(0)\right]}_{\mathrm{ATE}}$$

The curve reported in Figure A7 suggest that there is no significant treatment heterogeneity that can be predicted by our model of the CATEs.[38] This could be due to the limited statistical power of this analysis—that requires splitting the data in two halves—in our context with small treatment effects.

[38] We refer to the `grf` package online tutorial for further details on the construction of this figure.

FIGURE A7.—TOC ALONG PREDICTED CATE (Yadlowsky et al., 2021)

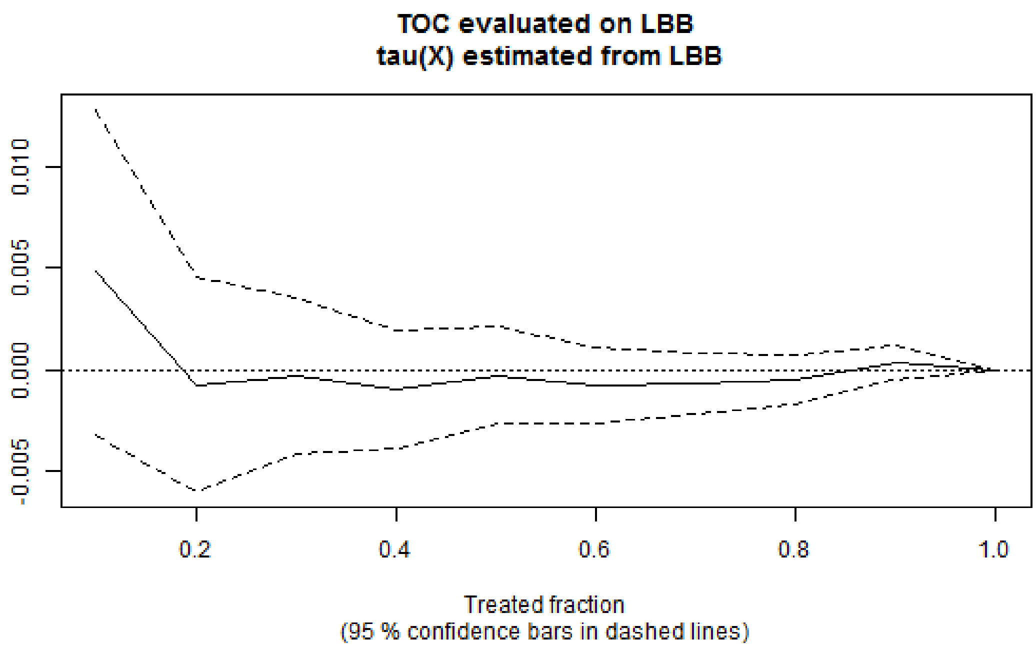

Note: This figure reports the TOC parameter (defined in the main text) from Yadlowsky et al. (2021) for $q \in (0,1]$. The solid line reports the TOC—that is the gap between the ATE among the fraction q of job seekers with the highest predicted CATEs—and the overall ATE from treating everyone. The predictive model of the CATEs was built on a random half of our sample, while the estimation of the TOC curve reported above was made on the second half. By construction, this curve equals 0 when $q = 1$. The dotted lines give the 95% confidence interval around this quantity. For any q, the graph suggests that the ATE among the fraction q of job seekers with the highest predicted CATEs is not significantly larger than the overall ATE from treating everyone. As in the main text, the outcome is the job finding in short term contracts.

The second approach we explored is the one proposed by Chernozhukov et al. (2018). This paper offers two tests of treatment effect heterogeneity, the Best Linear Predictor (BLP) test and one based on the sorted Grouped Average Treatment Effects (GATEs). Both tests conclude to no significant treatment effect heterogeneity in our data—we report in Figure A8 the results for the GATEs analysis. We refer to Chernozhukov et al. (2018) for further details on both tests (BLP and GATEs). We implemented those in R using the R package developed here.

FIGURE A8.—GROUPED AVERAGE TREATMENT EFFECTS (GATES) Chernozhukov et al. (2018)

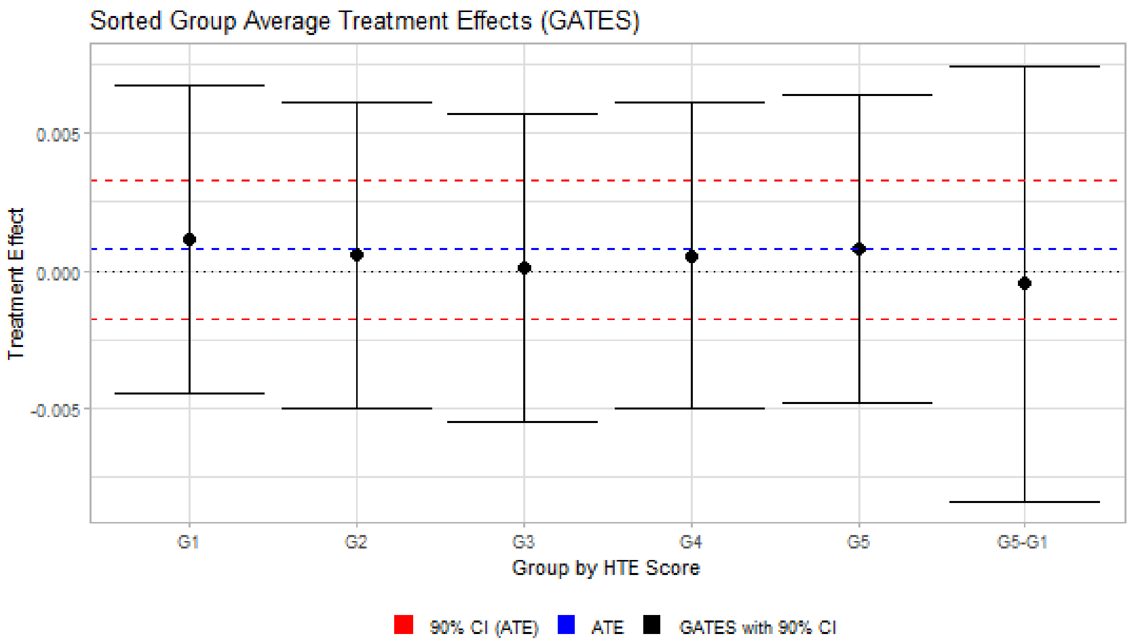

Note: This graph reports the estimates of the average treatment effect among groups defined based on a machine learning prediction of the CATEs—the first group being the one with the lowest predicted CATEs, and the fifth group being the one with the highest predicted CATEs. The graph shows than there is not any treatment effect heterogeneity detected between the first and the fifth group.

Classical heterogeneity analyses Focusing on the impact of our intervention on the job finding rate in short-term contracts, we observe some level of treatment effect heterogeneity in Table A3—that remains statistically significant after correcting for multiple hypothesis testing using Anderson sharpened q-values (Anderson, 2008).[39] The effect appears to be statistically significant and of a larger than average magnitude among female job seekers, and job seekers holding a high school diploma. On the contrary, age does not appear as highly predictive of any treatment effect heterogeneity. Lastly, job seekers with above-median unemployment duration (at the time of the launch of the experiment) also display a marginally significant average treatment effect, of a larger than average magnitude.

TABLE A3

EFFECT ON SHORT TERM JOB FINDING BY JOB SEEKER TYPE

	(1) Female	(2) Male	(3) No diploma	(4) Diploma	(5) Age $<$ 38	(6) Age $\geq$ 38	(7) Spell $<$ p(50)	(8) Spell $\geq$ p(50)
Treatment	0.0029	-0.0004	-0.0008	0.0028	0.0015	0.0016	0.0007	0.0022
	(0.0011)	(0.0013)	(0.0013)	(0.0011)	(0.0013)	(0.0011)	(0.0012)	(0.0011)
	[0.01]	[0.76]	[0.56]	[0.01]	[0.25]	[0.15]	[0.58]	[0.05]
	[[0.042]]	[[0.613]]	[[0.496]]	[[0.042]]	[[0.334]]	[[0.231]]	[[0.496]]	[[0.112]]
Observations	439,443	360,854	313,852	486,445	406,181	394,116	400,737	399,560
Baseline	0.14	0.17	0.15	0.16	0.19	0.12	0.17	0.14

Notes: Standard errors clustered at the labor market (CZ $\times$ Occ.) level reported in parenthesis. The original p-values are reported under brackets. Lastly, under double brackets are reported the Anderson sharpened q-values (Anderson, 2008), that is p-values adjusted to control for the False Discovery Rate (FDR) — the FDR being the expected proportion of all rejections that are type I errors.

A.8. *Estimation of targeting and residual effects*

Recall that the parameters we are interested in are the following. The residual effect is then defined as the increase in the likelihood that *any* match (job seeker i, hiring firm j)

[39] Anderson (2008) offers a way to correct p-values for multiple hypothesis testing in order to control the False Discovery Rate (FDR)—the FDR being the expected proportion of all rejections that are type I errors.

occurs when job seeker i is in the treated group, in the absence of any recommendation for the pair (i,j). Formally, if Y_i^j denotes the indicator for whether job seeker i was hired in firm j, Z_i indicates whether or not i is in the treated group, and R_i^j indicates whether the pair (i,j) has been recommended, the residual effect is defined as:

$$\text{RESIDUAL} \equiv E\left[Y_i^j(Z_i = 1, R_i^j = 0) - Y_i^j(Z_i = 0, R_i^j = 0)\right]$$

In the core of the paper, we focus on the residual effect on *recommended* matches, which is defined as follows:

$$\text{RESIDUAL}^{R=1} \equiv E\left[Y_i^j(Z_i = 1, R_i^j = 0) - Y_i^j(Z_i = 0, R_i^j = 0)|R_i^j = 1\right]$$

On the other hand, the targeting effect is defined as the impact on the likelihood that a given match (i,j) occurs if job seeker i is treated and firm j was recommended to it. Formally:

$$\text{TARGETING}^{R=1} \equiv E\left[Y_i^j(Z_i = 1, R_i^j = 1) - Y_i^j(Z_i = 1, R_i^j = 0) \mid R_i^j = 1\right]$$

All the above quantities can be estimated as follows. For the overall residual effect (RESIDUAL), we consider the sample of (i) all potential matches involving a control job seeker, and (ii) potential matches involving a treated job seekers *and* that were not recommended. Then, we estimate RESIDUAL as follows:

$$\text{RESIDUAL} = \frac{1}{|D_{(Z_i=1,R_i^j=0)}|} \sum_{(i,j)\in D_{(Z_i=1,R_i^j=0)}} \left\{\frac{Y_i^j}{1-\pi_{ij}}\right\} \quad - \quad \frac{1}{|D_{(Z_i=0)}|} \sum_{(i,j)\in D_{(Z_i=0)}} \left\{Y_i^j\right\}$$

where $D_{(Z_i=z,R_i^j=r)}$ is the set of potential matches where the job seeker has treatment status z, and the dyad (i,j) has recommendation status r. And we have used $\pi_{ij} \equiv \Pr[R_i^j =$

$1|X_i, X_j]$, which is the theoretical probability that dyad (i,j) is recommended conditional on the observable characteristics of the i and j. This can is given by our recommendation algorithm.

For the residual effect on the recommended matches (RESIDUAL$^{R=1}$), presented in the core of the paper, we take the same sample as for RESIDUAL but the re-weighting scheme changes. We draw fake recommendation statuses for potential matches involving control job seekers. We then estimate using the following estimator:

$$\begin{aligned}\widehat{\text{RESIDUAL}}^{R=1} =& \frac{1}{|D_{(Z_i=1,R_i^j=0)}|} \sum_{(i,j)\in D_{(Z_i=1,R_i^j=0)}} \left\{ \frac{1-\bar{p}_{Z=1}}{\bar{p}_{Z=1}} \frac{\pi_{ij}}{1-\pi_{ij}} Y_i^j \right\} \\ & - \frac{1}{|D_{(Z_i=0)}|} \sum_{(i,j)\in D_{(Z_i=0)}} \left\{ R_i^j Y_i^j + (1-R_i^j) \frac{1-\bar{p}_{Z=0}}{\bar{p}_{Z=0}} \frac{\pi_{ij}}{1-\pi_{ij}} Y_i^j \right\}\end{aligned}$$

where $\bar{p}_{Z=z}$ denotes the empirical probability that a given dyad is recommended among matches involving job seekers with treatment status $Z=z$.

Lastly, we estimate the targeting effect on the recommended matches, TARGETING, by taking the sample of recommended matches and computing the following estimator:

$$\begin{aligned}\widehat{\text{TARGETING}} =& \frac{1}{|D_{(Z_i=1,R_i^j=1)}|} \sum_{(i,j)\in D_{(Z_i=1,R_i^j=1)}} \left\{ Y_i^j \right\} \\ & - \frac{1}{|D_{(Z_i=1,R_i^j=0)}|} \sum_{(i,j)\in D_{(Z_i=1,R_i^j=0)}} \left\{ \frac{1-\bar{p}_{Z=1}}{\bar{p}_{Z=1}} \frac{\pi_{ij}}{1-\pi_{ij}} Y_i^j \right\}.\end{aligned}$$