

Encouraging and Directing Job Search: Direct and Spillover Effects in a Large Scale Experiment

Luc Behaghel, Sofia Dromundo, Marc Gurgand, Yagan
Hazard and Thomas Zuber¹

December 2022, WP #900

ABSTRACT

We analyze the employment effects of directing job seekers' applications towards establishments likely to recruit, building upon an existing Internet platform developed by the French public employment service. Our two-sided randomization design, with about 1.2 million job seekers and 100,000 establishments, allows us to measure precisely the effects of the recommender system at hand. Our randomized encouragement to use the system induces a 2% increase in job finding rates among women. This effect is due to an activation effect (increased search effort, stronger for women than men), but also to a targeting effect by which treated men and women were more likely to be hired by the firms that were specifically recommended to them. In a second step, we analyze whether these partial equilibrium effects translate into positive effects on aggregate employment. Drawing on the recent literature on the econometrics of interference effects, we estimate that by redirecting the search effort of some job seekers outside their initial job market, we reduced congestion in slack markets. Estimates suggest that this effect is only partly offset by the increased competition in initially tight markets, so that the intervention increases aggregate job finding rates.

Keywords: Search and Matching, Occupational Mobility, Displacement Effects


JEL classification: E24, J60, J62, J64

¹ Luc Behaghel: Paris School of Economics and INRAE. Sofia Dromundo: OECD. Marc Gurgand: Paris School of Economics and CNRS. Yagan Hazard: Paris School of Economics and Chaire Travail. Thomas Zuber: Banque de France. 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, PSE, Tinbergen Institute, and the 2020 "Matching jobs and workers online" IZA workshop. All remaining errors are ours.


NON-TECHNICAL SUMMARY

In this paper, we provide experimental evidence on the potential social value of a large-scale public recommendation algorithm developed by the French public employment service (PES). The platform is called “La Bonne Boîte” (“The adequate firm”, henceforth LBB). The platform started in 2015, and is based on an algorithm predicting hirings at the firm/occupation level. The goal of PES with this service is to provide job seekers with access to the so-called “hidden market” of firms that recruit without necessarily posting job ads. On the business-as-usual mode, the LBB website directs job seekers toward a list of firms most likely to hire them according to the location and occupation criteria they enter. We partner with PES to test the impact of this service using a randomized encouragement design: we send emails to about 800,000 registered job seekers (the treatment group) to encourage them to use LBB, and measure the impact on job finding rates. To analyze mechanisms and potential improvements, we use the encouragement emails sent to treated job seekers one step further, in the form of targeted recommendations toward specific firms within and outside their occupation of reference. While we introduce some random variation when making these recommendations, we also discipline ourselves using a simple, flexible equilibrium model at the commuting zone level. The model takes into account information on local tightness across occupations and makes educated guesses on key parameters (occupational mobility costs, firms' screening technology) to optimize recommendations in order to maximize the expected local aggregate employment. In a second step, we analyze ex post whether these presumed optimal recommendations were indeed effective. Specifically, we ask two questions: (i) Is there a positive private return to the email's recommendations and encouragement to search via LBB? This is directly identified by the reduced form effect comparing treated job seekers (who received the email) to control ones (who received no email). (ii) Do the recommendations generated by the ex-ante model strike the right balance, in terms of the breadth of occupational search, between congestion and mobility costs? Answering this second question is harder, as it involves estimating interference effects: recommendations made to a given job seeker, if they lead to a change in their application behavior, are likely to have external effects by displacing other job seekers. We build upon the recent literature on interference in randomized trials, in particular Hu, Li and Wager (2022), to estimate not only the direct effect but also the indirect effect of recommendations.

Graph 1a: Treatment effect on job finding rate for men.



Graph 1b: Treatment effect on job finding rate for women.



Note: These graphs report the treatment effect estimates for job finding at different time horizons for (a) men and (b) women. Standard errors are clustered at the labor market (Occ.*CZ) level and associated 95% confidence intervals are displayed graphically.

We find that the e-mails' recommendations and search encouragement increase by around 1.5% the job finding rate of female job seekers (+0.26pp from a baseline of 17.43%, see graph 1b). This result is driven by an activation effect of our intervention that led to an overall increase in search effort for women. Contrary to women, men appear to concentrate their search effort on the narrow set of LBB and recommended firms ("targeting effect", see graph 1a). This redirection of men's job search effort leads to a relative fall in job finding among non LBB firms and rise in job finding in LBB firms, resulting in an overall effect which is not statistically different from zero.

In a next step we document that the "targeting effect" (i.e. an increase in the likelihood that specific matches between pairs of job seekers and firms occur) is especially strong when we recommend firms hiring outside of a jobs seeker's origin occupation. This underlines the ability of the recommender system to redirect job seekers' search effort, leaving room for a potentially beneficial reallocation of job seekers' search effort across labor markets. We find that recommending job seekers to search in nearby occupations significantly reduces congestion frictions in their origin occupation. This is especially true when origin occupation were initially slack, i.e. when congestion frictions on the job seekers' side were initially high. Overall, this provides evidence, in a real set-up, that recommender systems can be used to reduce mismatch unemployment due to informational frictions.

Encourager et rediriger l'effort de recherche des demandeurs d'emploi : effets directs et indirects d'une expérience à grande échelle

RÉSUMÉ

Nous présentons dans cet article les résultats d'une expérience aléatoire à grande échelle consistant à mettre en avant auprès de demandeurs d'emploi des entreprises susceptibles de recruter dans un avenir proche. L'envoi aléatoire de conseils de candidature personnalisés induit une augmentation de 2 % du taux de retour à l'emploi chez les femmes. Cet effet est dû à un effet d'activation (augmentation de l'effort de recherche, plus fort pour les femmes que pour les hommes), mais aussi à un effet de ciblage des entreprises recommandées. Au niveau agrégé, nous montrons que la redirection de l'effort de recherche des demandeurs d'emploi vers des marchés en tension permet d'améliorer les appariements sur le marché du travail.

Mots-clés : chômage, appariements sur le marché du travail, effets de congestion, mobilité occupationnelle.

Les Documents de travail reflètent les idées personnelles de leurs auteurs et n'expriment pas nécessairement la position de la Banque de France. Ils sont disponibles sur publications.banque-france.fr

I 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 increase 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 to apply to specific firms to which they might not have spontaneously applied.

The rationale is that such services may increase equilibrium employment by reducing the cost of search frictions, through two distinct mechanisms: a reduction in individual job seekers' or firms' search cost, and a reduction in occupational mismatch if the platform's recommendations to job seekers are targeted toward tighter occupations. In practice, this simple theory of change may however be questioned for two related reasons. First, to reduce informational frictions, the PES' advice needs to be based on better information than privately available to job seekers and firms. If not, irrelevant recommendations may actually increase search cost as they push job seekers to apply where they would not be productive, in turn overwhelming firms with bad candidates and increasing their screening costs. In that respect, the PES' informational advantage likely lies in the administrative data it has on the supply and demand of labor at the local level. This allows to identify local mismatches that may not be visible to individual players. However, here comes the second difficulty: to leverage its informational advantage, the PES needs to trade off occupational mobility costs against congestion externalities. Indeed, recommending the job seekers to broaden their search toward a tighter nearby occupation may help by reducing congestion effects, but it also imposes a mobility cost (adjust their skills to a different job, or working in a second-choice occupation). How far one should recommend job seekers to go is not obvious. To sum up, in order to reduce frictional unemployment, public recommendation algorithms must find reliable operational solutions on two fronts: gather reliable local tightness measures, and parameterize the algorithm so as to reduce mismatch without inducing excessive mobility costs.

In this paper, we provide experimental evidence on the potential social value of a large-scale public recommendation algorithm developed by the French public employment service. The platform is called “La Bonne Boîte” (“The adequate firm,” henceforth LBB). It was started in

¹See in particular Horton (2017); Kuhn and Skuterud (2004); Kuhn and Mansour (2013); Kuhn (2014); Belot et al. (2018b,a) for studies of job search platform / recommender systems. Kircher (2020) provides a recent review of this literature.

2015, based on an algorithm predicting hirings at the firm $\times$ occupation level. The goal of the PES with this service is to provide job seekers with access to the so-called “hidden market” of firms that recruit without necessarily posting job ads. On the business-as-usual mode, the LBB website directs job seekers toward a list of firms most likely to hire them according to the location and occupation criteria they enter. We partner with the PES to test the impact of this service using a randomized encouragement design: we send emails to about 800,000 registered job seekers (the treatment group) to encourage them to use LBB, and measure the impact on job finding rates. To analyze mechanisms and potential improvements, we use the encouragement emails sent to treated job seekers one step further, in the form of targeted recommendations toward specific firms within and outside their occupation of reference. While we introduce some random variation when making these recommendations, we also discipline ourselves using a simple, flexible equilibrium model at the commuting zone level. The model takes into account information on local tightness across occupations and makes educated guesses on key parameters (occupational mobility costs, firms’ screening technology) to optimize recommendations in order to maximize the expected local aggregate employment. In a second step, we analyze *ex post* whether these presumed optimal recommendations were indeed effective. Specifically, we ask two questions : (i) Is there a positive private return to the email’s recommendations and encouragement to search via LBB? This is directly identified by the reduced form effect comparing treated job seekers (who received the email) to control ones (who received no email). (ii) Do the recommendations generated by the *ex ante* model strike the right balance, in terms of the breadth of occupational search, between congestion and mobility costs? Answering this second question is harder, as it involves estimating interference effects: recommendations made to a given job seeker, if they lead to a change in their application behavior, are likely to have external effects by displacing other job seekers. We build upon the recent literature on interference in randomized trials, in particular Hu et al. (2022), to estimate not only the direct effect but also the indirect effect of recommendations.

We find that the e-mails’ recommendations and search encouragement increase by around 1.5% the job finding rate of female job seekers (+0.26pp from a baseline of 17.43%). This seems to be primarily driven by an activation effect of our intervention that led to an increase in search effort. This additional effort appears to be concentrated on firms that are displayed on LBB’s online platform — no matter whether or not these firms were specifically recommended to the job seeker in the e-mail they received. Nevertheless, we also document a targeting effect of our intervention: an increase in the likelihood that specific matches between pairs of job seekers and firms occur when we recommend such matches in our emails. This underlines the ability of the recommender system to redirect search effort, leaving room for a potentially beneficial reallocation of labor across labor markets. As a last step of the analysis, we therefore document effects on aggregate employment. We estimate that recommending job seekers to search toward nearby tighter occupations significantly reduces congestion frictions in slack markets from which search effort was displaced. On the flip side, this increases frictions in the “destination markets”

of our recommendations. Yet the relative magnitude of these opposite effects suggest a net positive effect. Overall, this provides evidence, in a real set-up, that recommender systems can be used to reduce mismatch unemployment due to informational frictions.

This paper contributes to two main strands of literature. First, it fits into the growing literature on the labor market impacts of recommendation platforms. 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 paper that documented such effects in this fashion, highlighting the potential of tailored online screening of applicants to increase the vacancy filling rate on the firm side of the market. Rather concurrently, Belot et al. (2018b) documented the potential of customized online advice to broaden the occupational scope of some job search scope in a small-scale lab experiment. Since then, further work has started to study the optimal design of labor market recommender systems (Alfonso Naya et al., 2021). However, to the best of our knowledge our work is the first to document the effect of one such algorithm *at scale*, in an experiment involving roughly a quarter of the entire french labor market.

Second, our paper builds on, and provides a well-suited application to, the literature on the design and the evaluation of policy interventions in the presence of interference. The concern that the overall effect of interventions in various domains (labor, health or education, for instance) may differ substantively from their partial equilibrium effect is not new. A recent literature uses innovative experimental or quasi-experimental designs to quantitatively assess the underlying crowding-out (or crowding-in) effects. Seminal papers include Miguel and Kremer (2004) and Crépon et al. (2013). These papers use randomized saturation designs where some local markets have a lower proportion of treated units than others. Though particularly compelling, this approach builds on noisy comparisons across a small number of local markets, and therefore suffers from limited statistical power. In our case, the direct effects of the recommendations are at best small, so that indirect effects are very unlikely to be statistically detectable from such comparisons. As an alternative, Hu et al. (2022) have recently forcefully advocated the use of variations in indirect exposure to treatments within local markets. They introduce the “average direct effect” (ADE) and the “average indirect effect” (AIE) and show how the sum of these two effects directly translates into policy relevant parameters. Our paper provides an illustration of the value of their approach.

The paper proceeds as follows. Section II sets up a basic model to illustrate the trade-off between congestion and occupational mobility costs, and to derive a simple sufficient statistic to assess whether recommendation are too far or too close. In Section III, we provide background

information on LBB's job search platform. Section IV presents a workable solution to assign recommendations in an ex ante optimal way through emails advertising the platform. Section V presents the evaluation design and the data. Section VI provides the results on the private return to receiving the emails, and decomposes this impact into an activation and a targeting effect. Section VII analyzes whether our intervention generated social returns by reducing congestion frictions in slack markets, and reallocating labor to tight markets to help labor demand meet supply. Section VIII concludes.

II An illustrative model of optimal recommendations with interference effects

II.1 Identifying the direct and indirect effects of job search recommendations

Consider a local labor market with two homogeneous job seekers (indexed by $i \in \{1, 2\}$). They are looking for a job in the same occupation and the same commuting zone, or equivalently the same "local market." Two firms are willing to hire. The first firm, indexed by $j = 0$, looks for workers in the occupation of the job seekers: denoting by $d_{i,j}$ the "occupational distance" between firm j and the job seeker i , this implies $d_{i,0} = 0$ for both job seekers $i \in \{1, 2\}$. The second firm, indexed by $j = 1$, looks primarily for workers in a neighboring occupation, which we denote by $d_{i,1} = 1$ for $i \in \{1, 2\}$. Under the assumption that skills are not perfectly transferable from one occupation to the other, both job seekers are more likely to be hired when they apply to firm 0 in their own occupation. If both job seekers were to apply to the same firm, however, congestion or rationing effects may weigh negatively on each individual job seeker's labor market outcomes. Indeed, when labor demand at each firm is not perfectly elastic, each job seeker probability of being hired in a given firm will depend on his own as well the other job seeker's application behavior. In practice these congestion effects might entice a social planner to divert some job seekers away from firm 0 and toward firm 1 in order to increase the aggregate job finding rate of the economy.

Consider a policy intervention that sends one recommendation to each job seeker. The recommendation can be either to apply to firm 0 or to firm 1. Let $R_i^j = 1$ denote the fact that we recommend firm j to job seeker i and $R_i^j = 0$ otherwise. Because each job seeker only receives one recommendation we know that:

$$R_i^0 + R_i^1 = 1$$

While this notation is very general and will be useful later on in order to accommodate settings with many job seekers and firms, for the purpose of the present simple example we will concentrate on the effect of a specific "treatment", namely directing job seekers away from firm 0 and toward firm 1 which is hiring in the neighboring occupation. To this end we define a more specific dummy

variable W_i which will take the value 1 if job seeker i is directed toward firm 1 as opposed to firm 0:

$$W_i = R_i^1 = 1 - R_i^0$$

We assume that the two treatment variables W_1 and W_2 are independent and drawn from the same Bernoulli distribution with mean π .² In this setting the parameter π governs the degree of job seekers reallocation across labor markets and would ideally need to be chosen optimally: with what probability should policy makers recommend that job seekers broaden their occupational search to neighboring markets? Assume that the objective function of the policy maker is the aggregate employment rate in this economy. As the two job seekers are homogeneous, it is equal to the employment probability of any of them, $\mathbb{E}(Y_i)$, where Y_i is the indicator variable equal to 1 if job seeker i is hired by one of the two firms. We assume that each job seeker only applies to one firm, so that he or she can only receive an offer from one of the two firms. Denoting Y_i^j the indicator variable equal to 1 if i is hired by j , one has

$$Y_i = Y_i^0 + Y_i^1.$$

Whether job seeker i is ultimately hired by firm j depends on three elements: (i) the occupational distance $d_{i,j}$ separating job seeker i 's skills from firm j 's requirements; (ii) conditional on $d_{i,j}$, on whether i applied to j ; (iii) on whether the other job seeker ($-i$) also applied to j . At this point let us make no explicit assumption on job seekers' application behavior, except that application decisions are taken independently. We assume that the application behavior of i toward j only depends on W_i while the application behavior of the other worker $-i$ only depends on W_{-i} . In general we can define potential outcome variables for each job seeker/pair (i, j) as:

$$Y_i^j(R_i^j, R_{-i}^j).$$

Recalling that W_i stands for the fact of recommending firm 1 in the neighboring market as opposed to firm 0 in one's own market, potential outcomes can be re-written directly as functions of W_i and W_{-i} instead of the more general R_i^j . Under this convention worker i 's potential outcomes write

$$Y_i^0(1 - W_i, 1 - W_{-i})$$

in firm 0 and

$$Y_i^1(W_i, W_{-i})$$

in firm 1, with $(W_i, W_{-i}) \in \{0; 1\}^2$.

²This can be seen as from the researcher's perspective as a "Bernoulli trial" (with W_i being the treatment indicator), or as a "mixed strategy" from the policy maker's perspective. Of course, a corner solution may be optimal for the policy maker, with recommendation systematically made to firm 0 ($\pi = 0$) or to firm 1 ($\pi = 1$).

In this setting job seeker i 's probability of being hired in firms 0 and 1 are functions of π and of potential outcomes:

$$\mathbb{E}(Y_i^0) = (1 - \pi)^2 \mathbb{E}(Y_i^0(1, 1)) + (1 - \pi)\pi \mathbb{E}(Y_i^0(1, 0)) + \pi(1 - \pi) \mathbb{E}(Y_i^0(0, 1)) + \pi^2 \mathbb{E}(Y_i^0(0, 0))$$

and

$$\mathbb{E}(Y_i^1) = \pi^2 \mathbb{E}(Y_i^1(1, 1)) + \pi(1 - \pi) \mathbb{E}(Y_i^1(1, 0)) + \pi(1 - \pi) \mathbb{E}(Y_i^1(0, 1)) + (1 - \pi)^2 \mathbb{E}(Y_i^1(0, 0)).$$

Given that job seeker i 's overall job finding rate is just the sum of Y_i^0 and Y_i^1 , the effect of small change in π on job seeker i 's overall job finding rate can be expressed as a function of π and all eight potential outcomes as

$$\begin{aligned} \frac{\partial \mathbb{E}(Y_i)}{\partial \pi} = \frac{\partial \mathbb{E}(Y_i^0 + Y_i^1)}{\partial \pi} &= -2(1 - \pi) \mathbb{E}[Y_i^0(1, 1)] + (1 - 2\pi) \mathbb{E}[Y_i^0(1, 0)] \\ &\quad + (1 - 2\pi) \mathbb{E}[Y_i^0(0, 1)] + 2\pi \mathbb{E}[Y_i^0(0, 0)] \\ &\quad + 2\pi \mathbb{E}[Y_i^1(1, 1)] + (1 - 2\pi) \mathbb{E}[Y_i^1(1, 0)] \\ &\quad + (1 - 2\pi) \mathbb{E}[Y_i^1(0, 1)] - 2(1 - \pi) \mathbb{E}[Y_i^1(0, 0)]. \end{aligned} \quad (1)$$

Even though highly stylized, this model is rich enough to illustrate two important points. First, it shows that a policy maker can learn about the optimal π without necessarily testing alternative values of π across different local markets, as would be the case in a randomized saturation design. Instead, by picking a value of π (which can be chosen close to .5 in order to maximize statistical power, or close to priors on the optimal value π^*), the policy maker may estimate employment rates in firm j from a Bernoulli trial in which, due to random assignment, average potential outcomes are identified by

$$\mathbb{E}[Y_i^0(1 - W_i, 1 - W_{-i})] = \mathbb{E}[Y_i^0 \mid 1 - W_i, 1 - W_{-i}]$$

and

$$\mathbb{E}[Y_i^1(W_i, W_{-i})] = \mathbb{E}[Y_i^1 \mid W_i, W_{-i}].$$

With these potential outcomes in hand the social planner can compute the marginal effect on aggregate job finding of narrowing or widening job search and adjust π accordingly. The downside of this approach is of course that the policy maker only learns about the derivative of the objective function at the chosen value for π , $\partial \mathbb{E}(Y_i)/\partial \pi$, while a randomized saturation design varying π would identify $\mathbb{E}(Y_i \mid \pi)$ on a whole range of possible saturation levels (π).

Second, the expression in equation 1 is directly related to the sum of two average effects that Hu et al. (2022) call the average direct effect, τ_{ADE} , and the average indirect effect, τ_{AIE} . The average direct effect directly relates to what would be an average treatment effect in a design without interference: in our case the average effect of widening job seeker i 's occupational search on job seeker i 's own job finding rate. The average indirect effect on the contrary has no direct

counterpart in a setting without interference. In presence of cross treatments interference, the average indirect effect is the average effect that widening any job seeker's occupational search has on the aggregate job finding rate of all other job seekers. In this setting the total effect of treatment on the average outcome is the sum of the direct and indirect effects. Recommending job seekers to concentrate their job search effort on neighboring occupations may have a direct negative effect on job finding because of skill loss, but an indirect average positive effect through decreased competition in the origin occupation, hereby increasing the aggregate job finding rate.

In our simple example with just two firms and job seekers, these two effect have simple expressions. Defining worker i 's total potential outcome as a function of (W_i, W_{-i}) :

$$Y_i(W_i, W_{-i}) = Y_i^0(1 - W_i, 1 - W_{-i}) + Y_i^1(W_i, W_{-i})$$

we can apply the definition in Hu et al. (2022) to express the average direct effect as:

$$\tau_{ADE} = \mathbb{E}[Y_i(W_i = 1, W_{-i}) - Y_i(W_i = 0, W_{-i})],$$

and the average indirect effect as:

$$\tau_{AIE} = \mathbb{E}[Y_i(W_i, W_{-i} = 1) - Y_i(W_i, W_{-i} = 0)].$$

Given that both W_i and W_{-i} follow independent Bernoulli laws of intensity π and that job seeker i 's overall potential outcome Y_i is just the sum of Y_i^0 and Y_i^1 , the average direct and indirect effect are given by:

$$\begin{aligned} \tau_{ADE} = & (1 - \pi)\mathbb{E}[Y_i^0(1, 1) - Y_i^0(0, 1)] + \pi\mathbb{E}[Y_i^0(1, 0) - Y_i^0(0, 0)] \\ & + (1 - \pi)\mathbb{E}[Y_i^1(0, 0) - Y_i^1(1, 0)] + \pi\mathbb{E}[Y_i^1(0, 1) - Y_i^1(1, 1)] \end{aligned} \quad (2)$$

and

$$\begin{aligned} \tau_{AIE} = & (1 - \pi)\mathbb{E}[Y_i^0(1, 1) - Y_i^0(1, 0)] + \pi\mathbb{E}[Y_i^0(0, 1) - Y_i^0(0, 0)] \\ & + (1 - \pi)\mathbb{E}[Y_i^1(0, 0) - Y_i^1(0, 1)] + \pi\mathbb{E}[Y_i^1(1, 0) - Y_i^1(1, 1)] \end{aligned} \quad (3)$$

Comparing equations 1, 2 and 3, one can check that

$$\frac{\partial \mathbb{E}(Y_i)}{\partial \pi} = \tau_{ADE} + \tau_{AIE}. \quad (4)$$

Equation 4 is a direct application of Theorem 1 in Hu et al. (2022). It states that the marginal impact of increasing the probability to recommend firm 1 at the expense of firm 0 is equal to the sum of an average direct effect (τ_{ADE}) and an average indirect effect (τ_{AIE}).

II.2 Congestion effects versus occupational distance

Keeping the exact same setting with two job seekers and two firms, a “model” of the labor market consists in a set of values for all potential outcomes. Closely mimicking the model actually used

in our randomization process (see Section IV), let us assume that, conditional on receiving a recommendation, worker i applies with probability 1 at firm j whenever $d_{i,j} = 0$ and with probability $\rho_W < 1$ whenever $d_{i,j} = 1$. We assume that if not recommended to do so, job seekers do not apply to any firm. Once job seekers' applications have been sent to a particular firm they pass through the receiving firm's internal screening process. If a firm receives one application only, this unique application simply goes through to following steps of the hiring process. If a firm receives more than application, each application makes it through to later stages of the hiring process with probability $c \in [0.5; 1]$. The parameter c hence measures the degree of firm level congestion in the hiring process. If $c = 0.5$, only one application gets selected on average when two are received: there is full congestion. If $c = 1$, all applications are considered by the firm, there is no congestion. Once an application successfully passes through the screening step, firms decide to hire or reject the candidate based on his or her labor market skills. Mirroring job seekers' application decisions, we assume that firm j decides to hire interviewed candidate j with probability 1 if $d_{i,j} = 0$ and with probability $\rho_F < 1$ if $d_{i,j} = 1$.

Under these stylized assumptions the full set of potential outcomes can be described by:

$$\begin{aligned} \mathbb{E}(Y_i^0(0,0)) &= 0 & \mathbb{E}(Y_i^0(0,1)) &= 0 \\ \mathbb{E}(Y_i^0(1,0)) &= 1 & \mathbb{E}(Y_i^0(1,1)) &= c \end{aligned}$$

and

$$\begin{aligned} \mathbb{E}(Y_i^1(0,0)) &= 0 & \mathbb{E}(Y_i^1(0,1)) &= 0 \\ \mathbb{E}(Y_i^1(1,0)) &= \rho_W \rho_F & \mathbb{E}(Y_i^1(1,1)) &= \rho_W^2 \times c \times \rho_F + \rho_W(1 - \rho_W)\rho_F \\ & & &= \rho_W \rho_F - (1 - c)\rho_W^2 \rho_F. \end{aligned}$$

In the absence of congestion effects (i.e. $c = 1$) the application behavior of other workers do not affect one's own outcome:

$$\begin{aligned} \mathbb{E}(Y_i^0(1,0)) &= \mathbb{E}(Y_i^0(1,1)) = 1 \\ \mathbb{E}(Y_i^1(1,0)) &= \mathbb{E}(Y_i^1(1,1)) = \rho_W \rho_F. \end{aligned}$$

Moreover, the effective congestion effect defined as the percentage fall in hiring probabilities when one more worker gets recommended to the same firm decreases in absolute value from $c - 1 < 0$ at $d_{i,j} = 0$ to $\rho_W(c - 1)$ at $d_{i,j} = 1$. The decline of congestion effects with occupational distance arises because less and less recommendations transform into actual applications once $d > 0$.³

Substituting for potential outcomes in the marginal effect of π defined above we get:

$$\frac{\partial \mathbb{E}(Y_i)}{\partial \pi} = -2(1 - \pi)c + (1 - 2\pi) + 2\pi[\rho_W \rho_F - (1 - c)\rho_W^2 \rho_F] + (1 - 2\pi)\rho_W \rho_F$$

³We define the effective congestion effect at firm j as:

$$\frac{\mathbb{E}(Y_i^j(1,1)) - \mathbb{E}(Y_i^1(1,0))}{\mathbb{E}(Y_i^1(1,0))}.$$

Given this expression and assuming that $E(Y_i)$ is a concave function of π we can solve for π^* the optimal degree of redirection of workers away from their origin occupation — $\pi^* = 0$ corresponding to no redirection at all while $\pi^* > 0$ entails that some workers are recommended to firms hiring outside of their origin occupation.⁴

A solution to this problem displaying a positive amount of worker reallocation across labor markets exists if and only:

$$1 + \rho_W \rho_F > 2c.$$

As a consequence labor market reallocation will never happen when the cost of occupational distance is high (i.e. if ρ_F and/or ρ_W are sufficiently close to zero) and/or when congestion effects are low (c close to 1). Notice that if there is no cost to occupational switching whatsoever ($\rho_W = \rho_F = 1$) then $\pi^* = 1/2$ regardless of the degree of congestion effects c .

Overall, the fact that the optimal policy depends on mobility costs, congestion levels, and the elasticity of labor demand to labor supply, suggests that it is likely to vary across local labor markets defined by occupations and geographical location. The goal of this study is to explore empirically optimal reallocation policies in a real setting with thousands of firms and job seekers interacting in connected labor markets.

III Context: “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 them in their search by encouraging them to make unsolicited (spontaneous) applications. 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 the firm directly is given (see Figure A3).

The distinguishing feature of LBB is to recommend firms deemed likely to hire, whether they have posted a job advertisement or not. The rationale is to reduce informational frictions by allowing job seekers to apply to the “hidden job market” of firms that have potential vacancies that they fill without posting jobs (through internal referrals, for instance). To do so, LBB

⁴If an interior solution exists it is given by:

$$\pi^* = \frac{1 + \rho_W \rho_F - 2c}{2(1 - c)(1 + \rho_W^2 \rho_F)}.$$

uses administrative data covering the universe of French firms to derive hiring predictions at the establishment $\times$ occupation hiring predictions.⁵ LBB then defines for each occupation a specific predicted hiring threshold above which an establishment is deemed a “hiring firm” (sometimes denoted by BB for “Bonne Boîte” in french) for this specific occupation.⁶ If there is no such establishment, LBB’s search engine suggests to extend the search to a wider geographical area.

We do not have a leeway on the algorithm used to predict hiring, and take it as given. However, we are sufficiently confident in the quality of LBB’s prediction for our purpose: their prediction does explain realized hirings. Figure A6 plots the relationship between the log of firms’ average predicted hiring, within twenty equal-size groups, and the log of realized average hiring in each of those groups of firms. 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.

In its business-as-usual mode, LBB only recommends firms likely to offer jobs in the occupation the job seeker entered in the search engine. In the next section, we present the additional algorithm that we develop to generate recommendations to broaden the occupational search. Then, in Section V, we present the evaluation design to assess the impact of this experimental development.

IV A workable optimal recommendation system

While the simple model of section II underscores the fact that optimal recommendations tightly depend on the occupational structure of the labor market as well as the strength of potential congestion effects, the model is not general enough to generate actual labor market recommendations. To that end we extend the model of section II to a setting involving many workers and firms with occupational distances strictly greater than 1 while allowing for congestion effects to depend continuously on the number of applications received by each firm. Our experimental design will heavily rely on this more general model to generate recommendations that are not purely arbitrary, while introducing controlled sources of variation (see details in Section V below).

⁵These 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/>).

⁶As a consequence, a given establishment can be considered as a “hiring firm” for one occupation but not for another.

IV.1 General setting and notations

Let the economy be composed of a spatially homogeneous labor market populated by N job seekers and F firms. Each job seeker looks for a job in his origin as well as neighboring occupations. Each firm may recruit workers in different occupations. We assume that hiring decisions are taken at the firm/occupation level and are not correlated within firm across occupations. The purpose of our model is to understand the effect of recommending specific firms/occupations pairs to job seekers on realized worker/firm matches. The timing of actions is the following:

1. We recommend firms to workers.
2. Workers who are more or less averse to occupational switching choose or not to apply to these firms according to occupational distance.
3. Firms skim through the applications they receive and randomly decide to look more deeply into some of them.
4. Firms are more or less efficient at screening applications. More efficient firms will be able to review a greater number of applications.
5. Firms, which are more or less averse to occupational switching, review selected applications and decide whether or not to hire each reviewed applicant according to occupational distance.

To describe more formally the model of the labor market used to solve our matching problem, we introduce the following notations. Let w and f index individual workers and firms, $i, j, \dots$ index occupations (each worker has a single occupation while each firm operates in several), $d_{i,j}$ index the occupational distance between two occupations i, j , $V^{f,j}$ denote the vacancies at firm f in occupation j , m_f denote the efficiency of firm f 's screening technology, $\rho_f \in (0, 1)$ denote the occupational discount factor of firm f , $\rho_w \in (0, 1)$ denote the occupational discount factor of worker w and $T(w)$ the ex-ante number of recommendations that will be sent to worker w .

IV.2 Computing the expected number of worker/firm matches

Given the structural parameters of the model, the goal of a central planner is to maximize the expected number of worker/firm matches in the economy. The central planner's choice variable is the full distribution of possible worker/firm pairwise recommendations. In this section we first derive the planner's objective.

Let us consider a worker w whose occupation is i . The hiring process unfolds as follows. The central planner draws $T(w)$ recommendations for worker w according to the generalized Bernoulli distribution $\alpha_w = \{\alpha_w^{f,j}\}^{f,j}$. The number of recommendations $T(w)$ received by worker w is given

ex-ante. The total number of recommendations sent by the social planner is given by:

$$T = \sum_{w \leq N} T(w).$$

The probability that worker w is recommended to occupation j in firm f at least once in one of the $T(w)$ draws is:

$$P(R_w^{f,j} = 1) = 1 - (1 - \alpha_w^{f,j})^{T(w)},$$

where $R_w^{f,j}$ denotes stands for the recommendation dummy. We denote by $R^{f,j}$ the set of workers to which we recommend the firm/occupation pair (f, j) at least once:

$$R^{f,j} = \{w | R_w^{f,j} = 1\}.$$

Given that worker w has been recommended to (f, j) , he actually applies to (f, j) with probability:

$$P(A_w^{f,j} = 1 | w \in R^{f,j}) = \rho_w^{d_{i,j}}$$

where $A_w^{f,j}$ is an application dummy and $A^{f,j}$ is the set of applicants to (f, j) :

$$A^{f,j} = \{w | A_w^{f,j} = 1\}.$$

We assume that workers can only apply if recommended to do so, so that:

$$A^{f,j} \subset R^{f,j}$$

Hence, the number of workers who apply to (f, j) is:

$$W^{f,j} = \sum_w \rho_w^{d_{i,j}} (w \in R^{f,j})$$

In the tradition of statistical inference, let us take the average. The unconditional expectation of $W^{f,j}$ is given by:

$$E[W^{f,j}] = \sum_w \rho_w^{d_{i,j}} [1 - (1 - \alpha_w^{f,j})^{T(w)}],$$

and its variance is:

$$V[W^{f,j}] = \sum_w \rho_w^{d_{i,j}} [1 - (1 - \alpha_w^{f,j})^{T(w)}] [1 - \rho_w^{d_{i,j}} [1 - (1 - \alpha_w^{f,j})^{T(w)}]].$$

Given that (f, j) receives $W^{f,j}$ applications, it randomly selects among them the ones that will be considered for employment. This selection occurs through the firm specific screening technology q_f which takes as its only argument the branch specific slackness ratio $\theta^{f,j} = W^{f,j}/V^{f,j}$. We assume that q_f is non-increasing, that $q_f(0) = 1$, and that $q_f(+\infty) = 0$.⁷ Conditional on applying to (f, j) a worker has probability $q^{f,j}$ to be interviewed:

$$q^{f,j} = E[q_f(\theta^{f,j})]$$

⁷See appendix A.7 for further details.

This expectation can be approximated by:

$$q^{f,j} = q_f(E[\theta^{f,j}]) + \frac{V[\theta^{f,j}]}{2} \frac{\partial^2 q_f}{\partial \theta^2}(E[\theta^{f,j}]) + o(E[(\theta^{f,j} - E[\theta^{f,j}])^3])$$

Once (f, j) has selected the $q^{f,j} W^{f,j}$ workers it will interview, the probability of each of them to be hired is simply $\rho_f^{d_{i,j}}$. Coming back to worker w , its unconditional probability of being hired by (f, j) is:

$$P(w \in H^{f,j}) = \rho_f^{d_{i,j}} \times q^{f,j} \times \rho_w^{d_{i,j}} \times [1 - (1 - \alpha_w^{f,j})^{T(w)}]$$

Hence, ignoring the possibility of a worker being hired by two firms, the probability that worker w will be hired by some firm is:

$$P(w \in H) = 1 - \prod_{f,j} [1 - P(w \in H^{f,j})].$$

This can be approximated by:

$$P(w \in H) \sim \sum_{f,j} P(w \in H^{f,j})$$

i.e.

$$P(w \in H) \sim \sum_{f,j} \rho_f^{d_{i,j}} \times q^{f,j} \times \rho_w^{d_{i,j}} [1 - (1 - \alpha_w^{f,j})^{T(w)}].$$

If we follow these steps for each worker we find that the expected total number of hires in the economy can be written as:

$$M = \sum_{w,f,j} \rho_f^{d_{i,j}} \times q^{f,j} \times \rho_w^{d_{i,j}} \times [1 - (1 - \alpha_w^{f,j})^{T(w)}]$$

The problem of the central planner is to maximize M subject to:

$$\begin{aligned} \forall w, \sum_{f,j} \alpha_w^{f,j} &= 1 \\ \forall (w, f, j), 0 &\leq \alpha_w^{f,j} \leq 1 \end{aligned}$$

This problem has dimensionality $\#(\text{Workers}) \times \#(\text{Firms}) \times \#(\text{Occupations})$, which in practice is too large. To reduce the dimensionality of the problem we parameterize $\alpha_w^{f,j}$ using available information on workers and firms. Denote $X_{w,f,j}$ the vector of worker/firm/branch characteristics that will be used to predict $\alpha_w^{f,j}$. We assume that:

$$\alpha_w^{f,j} = \frac{\exp(X'_{w,f,j} \beta)}{\sum_{f,j} \exp(X'_{w,f,j} \beta)};$$

Hence the dimensionality of the problem is reduced to $\#(\text{worker/firm characteristics})$ so that, in the end, the maximization problem reduces to:

$$\max_{\beta} \sum_{w,f,j} \rho_f^{d_{i,j}} \times q^{f,j} \times \rho_w^{d_{i,j}} \times [1 - (1 - \frac{\exp(X'_{w,f,j} \beta)}{\sum_{f,j} \exp(X'_{w,f,j} \beta)})^{T(w)}].$$

In practice the vector X may include observable market level, worker level and firm level characteristics, taken both from observed data (firms vacancies, worker/firm occupational distance) and structural parameters of the model (ρ_f , ρ_w , the shape parameters of q_f .) The weight given by the optimal parameter β to these different components will depend on the occupational distribution of job seekers and firms within each geographically defined labor market. In a 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. It would be the exact opposite if worker and firm were to be evenly distributed across the occupational space.

V Experimental design

This study builds upon a two-sided randomization that creates random exposure to recommendations by LBB on the firms' and job seekers' sides. In this section we first describe how we selected treated firms and job seekers who were included in the experiment and then turn to the drawing of pairwise recommendations linking the two sides of the markets.⁸

V.1 Drawing treated job seekers and treated firms

All experimental treatments are assigned within commuting zones.⁹ Our experimental sample covers 94 out of the 404 French commuting zones,¹⁰ representing a pool of 1,209,859 job seekers and 98,366 hiring establishments.

The basic experimental treatment consists in increasing treated firms' and treated job seekers' exposure to LBB's job search services. First, we randomly select a subset of firms among those short-listed by LBB's algorithm. We stratify the random selection of treated firms within 5-digits sectors and above median/below median predicted hiring bins. During four weeks, selected "treated" firms are displayed in priority in response to job seekers' requests on the website, while the remaining "control" firms are not displayed (or displayed at the bottom of the list if there are too few treated firms satisfying the search criteria). Second, we randomly draw two thirds of the 1.2 million job seekers to receive two or four emails pushing the LBB service, with specific recommendations toward up to eight of the treated firms. We stratify the random selec-

⁸We do not insist on the firm-level randomization, whose analysis is the focus of a companion paper.

⁹When assigning treatment within a commuting zone, we do not distinguish across job seeker and establishment pairs by their geographical distance. Indeed, the existing evidence suggests that spatial mismatch is second order compared to occupational mismatch (Marinescu and Rathelot, 2018). The role of geographical distance can however be analyzed ex post based on remaining non-experimental variation; this is kept for further analysis.

¹⁰We randomly selected these 94 Commuting Zones out of all the 404 possible commuting zones. We stratified this random selection of treated commuting zones within tightness and size quintiles. For more details on Commuting Zones and local labor markets see Appendix Section A.3.

tion of treated job seekers within desired occupations and above median/below median bins of a linearly predicted exit rate out of unemployment.

We randomly draw 806,437 treated job seekers and 38,810 treated establishments. Because a large share of job seekers exited the unemployment pool in the short period separating randomization from the actual start of our experiment, we will ex-post restrict our analysis to the 533,557 treated and 266,740 control job seekers who were still registered with PES and had not found a job as of 19th november 2019.¹¹

The balance of job seekers' observable variables across treatment and control groups is presented in Table 1, keeping only job seekers still unemployed at the beginning of the experiment. Furthermore this table presents the p-values associated to an F-Test of the regressions of each observable on four indicator variables corresponding to the four job seekers' treatment arms.

Table 1: BALANCE TABLE FOR JOB SEEKERS IN TREATED CZ.

Variable	(1) C	(2)	(3) T	(4)	(5) T-C	(6)
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. Column (7) presents the F-Test p-values for the regressions the variable listed in the first column on four indicator variables corresponding to the four job seekers' treatment arms.

V.2 Drawing pairwise recommendations

V.2.1 Additional treatment arms

Even though the random selection of a pool of treated job seekers and a pool of treated establishments tells us which job seekers and which establishments will enter our pairwise recommendations, it does not tell us which specific pairwise recommendations will be formed. In particular, should a given job seeker receive recommendations only toward firm likely to hire in their local labor market, or should recommendations be broadened toward firms likely to hire in neighboring

¹¹This pre-treatment attrition rate is well balanced across treatment and control groups.

occupations? As illustrated in Section II and further detailed in Section IV, the optimal solution depends on the position of the job seeker in the occupational space, and on the relative tightness in the occupations that surround him. Importantly, the specific solution depends on unknown parameters on mobility cost and congestion effects. The goal of our experiment is precisely to learn about the optimal solution. This requires to introduce random variation in the recommendations made. We do so by hypothesizing several plausible parameter values related congestion and mobility costs, and randomly assigning those values across firms and workers, thus defining additional treatment arms among treated job seekers and treated firms.

Specifically, we do not a priori known (a) *how many* recommendations job seekers and establishments should receive for these recommendations to have an effect. Secondly, we do not a priori know (b) *how far* in the occupational space we should advise job seekers and establishments to look for jobs and employees. In order to get a sense for (a) and (b) we build into our experimental design a further level of randomness by distributing 4 possible treatment status among treated job seekers and establishments, using a factorial design. Hence while among treated job seekers some will receive **many** recommendations, others will only receive a **few**. At the same time some treated job seekers will be recommended to establishments hiring **far away** in the occupational space while others will be recommended to establishments hiring **close to** their own occupation. Similarly, while some establishments will be recommended to large pool of job seekers *conditional* on their level of predicted hiring some other establishments will only be recommended to few job seekers. And while some establishments will be recommended to occupationally close-by job seekers, others will be recommended to job seekers far away in the occupational space. We sum up the structure of our experimental design and the distribution of the different treatment status for job seekers and establishments in Table 2.


Table 2: TREATMENT ARMS AND RECOMMENDATIONS TYPES

Job-seekers				Establishments			
Treated			Control	Treated			Control
	<i>Few</i>	<i>Many</i>			<i>Few</i>	<i>Many</i>	
<i>Close</i>	133,558	133,619	266,740	<i>Close</i>	9,716	9,614	59,556
<i>Far</i>	133,169	133,411		<i>Far</i>	9,792	9,688	

V.2.2 Applying the optimal recommendation algorithm

Based on their treatment arm, we assign job seekers and firms with specific values of the key parameters of the model of Section IV. The first one is $T(w)$ the number of recommendations received by job seeker w , which we take to be four in the “few” treatment arm and eight in

Figure 1: TREATMENT ARMS FOR JOB SEEKERS AND FIRMS



the “many” treatment arm.¹² The other parameters are ρ_w and ρ_f , the occupational distaste parameters on the worker’s and firm side respectively. According to each agent’s “far” or “close” treatment arm we select either a high value of these count factor, corresponding to a low distaste for occupational distance in the “far” group, or a low value corresponding to a strong distaste for occupational distance in the “close” group. Finally we model the strength of firm level congestion effects through a firm specific shift term m_f entering the screening function q . Firms in the “many” treatment arm are characterized by a high value of the screening efficiency m_f while firms in the “few” treatment arm are attributed a low value of m_f . As a consequence our recommendation algorithm should attribute relatively more recommendations to the “many”-type firms than to the “few”-type ones.¹³

With these random structural parameters in hand we turn to the recommendation model described in Section IV. We take firm/occupation level predicted hirings as our empirical counterpart of opened vacancies and solve for the optimal weights β in each of the 94 commuting zone. As could be expected, occupation distance as well as agents’ distaste of it affect the probability of a far away recommendation negatively. The firm level screening efficiency parameter attributed in the “few”/“many” treatment arms increases the expected number of recommended job seekers. Finally, everything else equal, large firms are also more likely to receive many recommendations. Once the optimal weights β are numerically solved for in each commuting zone we proceed to

¹²In practice, a job seeker assigned to the **many** treatment may not receive eight distinct recommendations, if the same firm/occupation pair is drawn more than once.

¹³In practice we select $T(w) \in \{4, 8\}$, $\rho_w \in \{0.82, 0.94\}$, $\rho_f \in \{0.82, 0.94\}$ and $m_f \in \{0.5, 1.5\}$. The curvature of the screening function is set to 3, see appendix A.7 for further details.

draw as many job seeker/firm/occupation recommendations as needed following the generalized Bernoulli distribution described in section IV.

For instance, for a job seeker i assigned to the “many” and “far away” treatment arms, we draw eight times from the pool of occupation/firm pairs indexed by (f, j) with a probability $\alpha_i^{f,j}$ where $\alpha_i^{f,j}$ is the optimal solution for a job seeker in this particular local market who was attributed a low mobility cost (large ρ_w), given the mobility costs, screening efficiencies and predicted hirings of all job seekers and firms who surround him or her.

In the end, on both sides of the market, each agent’s treatment status determines how many recommendations he will receive and how far these recommendations will be in the occupational space. Hence, while our pairwise recommendations partly reflect the non-random empirical distribution of job seekers and predicted vacancies across the occupational space, they also incorporate a random component linked to each agent’s specific treatment status which will allow us to identify the effect of the number of recommendations and their occupational distance.

As can be seen in Table 3, on average job seekers belonging to the “Few” treatment arm received recommendations to 3.19 distinct establishments while job seekers belonging to the “Many” treatment arm, received recommendations to 5.62 distinct establishments. In both the “Few” and “Many” treatment arms, the relative occupational distance of these recommendations varied according to each job seeker’s “Close” or “Far” treatment status. Whereas job seekers bound to receive “Close” recommendations were kept at a 0.55 average distance, job seekers in the “Far” treatment arm were set recommendations on average 1.28 occupations away from their original search occupation.

Table 3: JOB SEEKERS’ REALIZED TREATMENT

Variable	Group	Mean	Sd	Min	Max	Obs
Distinct rec.	Few	3.19	1.07	1	4	399821
	Many	5.62	2.34	1	8	399938
Occupational dist.	Close	0.55	1.19	0	15	400504
	Far	1.28	1.56	0	15	399705

Note: This table gives descriptive statistics for the number of distinct recommended firms in the “Few” versus “Many” job seekers’ treatment arms as well as the average occupational distance of job seekers’ recommended establishments in the “Close” versus “Far” treatment arms.

V.3 Emailing the job seekers

In practice, our experiment consists in emailing treated job seekers with links to LBB’s contact information of specific establishments. Job seekers interested in the establishment that we recommended can thus contact the firm and make an unsolicited application. Importantly the contact

information usually consists of a location, an email or a telephone number. When no contact information is available for a given establishment, LBB allows its user to directly search for this information on Google. Moreover, in some cases LBB allows job seekers visiting its pages to directly send an application through public employment services' online application tool. When this tool is available, and as can be seen in Figure A3 in appendix, job seekers just need to click on a "Send an application" (in French "Postuler") icon which appears on the right hand side of the contact information page.

As can be seen in Table 4 below or Figure A4 in appendix, the emails we used to direct job seekers to specific establishments contained the following information: the job seeker's name, general information on the hiring behavior of firms - and in particular on the fact that a considerable share of hirings stem from unsolicited applications -, general information on LBB, each job seekers desired occupation, at most two links to the LBB page of recommended establishments and, finally, a general purpose link directing toward LBB's search engine. Apart from the job seeker's name and search occupation the only specifically individual content of these emails were the links to the contact information of recommended firms. Importantly these links were job seeker/establishment specific so that by tracking job seekers' clicks we could record their interest in some specific establishment. How were this links formed and dispatched into different emails? As previously explained we drew within the pool of nearby treated establishments as many establishments, i.e. either 4 or 8, as each job seeker's treatment status required. Once these 4 or 8 recommendations had been drawn for each job seeker we distributed them respectively into either 2 or 4 different emails. Each email thus contained at most two links directing to the contact information of at most two distinct establishments. When a single establishment ended up appearing twice in a single email we collapsed the two links into one single link. Finally we distinguished between establishments hiring in a job seeker's own occupation and establishments hiring in another occupation by explicitly acknowledging one of the two cases when introducing each link. Establishments hiring in one's own occupation were introduced as such while establishments hiring in a neighboring occupation were framed as "hiring in an occupation not far from yours". After the specific links to recommended establishments' contact information, the email concluded with a general purpose link directing to LBB's search engine. The content of our emails is summed up in Table 4 below.

Table 4: AN EMAIL’S SCHEMATIC CONTENT

Dear Mr./Mrs. [X],

You are currently registered with the public employment services and are looking for a job as a [X’s occupation].

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 public employment services, has selected for you a few 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,

V.4 Reallocating labor across tight and slack markets

As we already made it clear in previous sections, our experiment aims at uncovering the potential of recommender systems to reduce mismatch unemployment. Indeed, this appears to be the main channel through which such devices could generate social returns.

A necessary condition for such recommender system to reduce mismatch unemployment is that it generates recommendations from slack labor markets (where labor is too abundant compared to the amount of posted vacancies) to tight markets. In that way, one can hope that the system reduces congestion frictions in slack markets while helping labor demand to meet supply in tight ones. Table 5 checks that it is the case of our recommendation algorithm in practice.

The average market in our experiment has around 23 job seekers for 10 hiring firms (BB). Its tightness, defined as the number of predicted hiring in hiring firms over the number of job seekers, is at 0.37. The median number of neighboring markets that are directly connected to any given market m in the occupational graph we use is 2. However, this hides a lot of heterogeneity across markets. This is why we distinguish in panels B and C of Table 5 two kinds of markets — the so-called “source” and “destination” markets. Source markets are those for which our algorithm generated an above median probability to make recommendations to neighboring occupations. In other words, these are the markets from which we re-oriented job seekers the most. These markets differ from the average market in key dimensions. They are smaller on average, with around 10

job seekers for 3 hiring firms. More importantly, their tightness measure is below the average over all markets, at 0.17. As mentioned above, this is a key property to be fulfilled for our algorithm to generate social gains: it should reallocate labor from slack to tight markets. Reassuringly, our algorithm appears to satisfy this condition, as our source markets seem to be slack markets. On the contrary, the destination markets — defined as those for which our algorithm generated an above median probability to make recommendations to neighboring occupations — appear to be bigger and tighter markets. The average number of job seekers in those markets is around 35, for 16 hiring firms. And their average tightness is at 0.56 (compared to 0.17 in source markets), suggesting that our redirection intervention towards such markets could contribute at reducing some existing mismatch unemployment.

Table 5: TIGHT AND SLACK LABOR MARKETS DESCRIPTIVE STATISTICS

Nb. markets	Nb. job seekers	Nb. hiring firms (BB)	Tightness	Median nb. of neighboring markets (d=1)
A. All markets				
35187	22.74 (72.82)	9.53 (37.05)	0.37 (1.01)	2.01 (1.66)
B. “Source” markets (above median prob. to recommend to neighboring occ.)				
17593	10.32 (45.16)	2.69 (18.66)	0.17 (0.61)	1.90 (1.69)
C. “Destination” markets (below median prob. to recommend to neighboring occ.)				
17594	35.17 (89.86)	16.37 (47.99)	0.56 (1.26)	2.13 (1.62)

Notes: Columns 2, 3 and 4 of this table report the average number of job seekers and hiring firms (BB) and the average tightness (as measured by the number of predicted hirings over the number of job seekers) for all or different categories of markets — with standard deviations reported in parenthesis. The last column reports the average number of directly neighboring markets — i.e., markets at distance 1 in the occupational graph we use. “Source markets” are defined as those within which job seekers faced a below median average probability to get recommendations outside of their market — the median probability being at 0.97. “Destination markets” are defined as those within which job seekers faced an above median average probability to get recommendations outside of their market.

VI Private returns to the encouragement: Activation and targeting effects

In this section, we start by providing reduced-form evidence that receiving the email increased job finding rates especially among female treated job seekers. We further decompose the reduced-form effects into a targeting and an activation effect, showing that the increased job finding rates occur through hirings both in recommended and in non-recommended firms. In a second step, we focus the analysis on pairs (dyads) consisting of a job seeker and a firm. We show that our design allows us to quantify the relative magnitude of the activation and the targeting effects, and confirm the existence of a stronger activation effect among female job seekers.

VI.1 Impact on job finding rates

We observe access to employment as registered by PES, over a period of four months since treatment. More specifically we know each job seeker's return to employment status, type of contract, the date at which this contract is set to start and, for definite duration contracts, the date at which this contract will be terminated. The main equation we estimate by OLS is the following:

$$y_i = \alpha + \beta Z_i + \epsilon_i.$$


The dependent variable of interest y_i corresponds to the job finding status of job seeker i at a given point in time, possibly conditioning on the type of contract found (finite or indefinite duration). Z_i is a dummy equal to 1 if job seeker i received an email. This model is estimated in the sample of treated and control individuals from treated commuting zones, and $\hat{\beta}$ is therefore an estimator of our "intention to treat" (ITT) parameter on the job seeker's side.¹⁴

Figure 2 presents the estimates of this ITT parameter at different time horizons pooling together all types of contract, for the subsample of male (panel a) or female (panel b) job seekers.¹⁵ Each point depicts the result of a separate regression of access to employment before some date on the treatment status Z_i for the set of job seekers who were still unemployed when our intervention began. Going from left to right, the time horizon widens so that the overall graph depicts the cumulative effect of our treatment on job finding.

¹⁴We describe β as an ITT parameter as a large share of "treated" job seekers ($Z_i = 1$) did not even open the e-mail we sent them in the first place. Therefore, we see our e-mail as an encouragement rather than a proper treatment in itself.

¹⁵Figure A8 in appendix reports the point estimates for the full sample, pooling male and female job seekers together.

Figure 2: JOB-FINDING RATE ITT ESTIMATES BY GENDER



Note: ITT estimates for job finding at different time horizons for (a) males and (b) females. Sample restricted to job seekers who were still unemployed as of 19/11/2019. Standard errors are clustered at the labor market (Occ.*CZ) level and associated 95% confidence intervals are displayed.

The respective responses of males and females to our intervention differ markedly. In Figure 2, the overall response of men is around zero, while women's response after two months since the beginning of our intervention is positive and significant. In Section A.6 in the appendix, we investigate candidate explanations for such heterogeneity — e.g., differences in observable characteristics and differences in take-up rates. Overall, this differential effect between males and females does not seem to be driven by differences in observables. However, female job seekers seem to respond more to our encouragement — in the sense that they open our e-mails more frequently than their males counterparts.¹⁶

Driven by this observation, we investigate whether the effect is concentrated on hiring firms (BB) that were recommended to job seekers. If so, this would suggest that the impact is mainly mediated by a reorientation of search effort to firms with better hiring prospects. Conversely, if we observe an effect of similar magnitude on job finding rates in firms that were not recommended to the job seeker, or in firms that were not even involved in the experiment because of their lower predicted hiring prospects, then we could conjecture that our impact is mainly driven by an activation effect. Table 6 reports the results of this decomposition exercise. There is evidence of an activation effect. Consider first panel A, that pools all job seekers. Looking at the last two columns that distinguish job finding in recommended and non-recommended hiring firms (BB), we observe that the effect among non-recommended firms (Non-rec. BB) is of the same order of magnitude as the one on recommended firms (Rec. BB) once related to their respected baseline — +14.7% of the baseline for recommended BB, +17.6% for non-recommended BB. The picture

¹⁶We also study the heterogeneity of our effect by contract types. Figure A9 in appendix A.6 shows that the effect on the job finding rate of female job seekers appear to be driven by an increase in the probability to find definite duration contracts.

stays quite similar when focusing on the subsample of female job seekers.¹⁷

¹⁷For male job seekers, the null effect seems to be driven by a decrease in their job finding rate in firms not presented on LBB (Not BB), that is counterbalanced by an increase in their job finding rate in hiring firms (BB). This increase is not particularly driven by an effect on the job finding rate in recommended firms — if anything, it is more pronounced in non-recommended BB.

Table 6: ITT ESTIMATES, BY TYPE OF FIRM

	All	Not BB	BB	Rec. BB	Non-rec. BB
A. All					
Baseline	0.19503 (0.00189)	0.18689 (0.00183)	0.00814 (0.00030)	0.00177 (0.00010)	0.00637 (0.00026)
Treatment	0.00070 (0.00092)	-0.00068 (0.00091)	0.00138 (0.00022)	0.00026 (0.00010)	0.00112 (0.00020)
Observations	800297	800297	800297	800297	800297
B. Females					
Baseline	0.17432 (0.00250)	0.16444 (0.00238)	0.00988 (0.00040)	0.00200 (0.00013)	0.00788 (0.00036)
Treatment	0.00261 (0.00119)	0.00127 (0.00115)	0.00135 (0.00033)	0.00034 (0.00015)	0.00100 (0.00029)
Observations	439443	439443	439443	439443	439443
C. Males					
Baseline	0.22033 (0.00181)	0.21431 (0.00174)	0.00602 (0.00035)	0.00150 (0.00014)	0.00452 (0.00028)
Treatment	-0.00173 (0.00145)	-0.00317 (0.00144)	0.00144 (0.00029)	0.00017 (0.00014)	0.00127 (0.00025)
Observations	360854	360854	360854	360854	360854

Notes: Standard errors clustered as the labor market (CZ $\times$ Occ.) level reported in parenthesis. Job finding rates displayed in the first column are decomposed into different categories of hiring, depending on the type of firm that made the hire. The coefficients displayed in the “Treatment” row are our ITT estimates, by type of firm. Column 1 reports the ITT estimate for all hiring. Column 2 (Not BB) focuses on firms that are not considered as “hiring firms” (i.e., their predicted hirings are not high enough according to LBB’s algorithm). Column 3 (BB) focuses on hirings firms (BB), whether they were recommended or not to the job seeker. Column 4 (Rec. BB) focuses on hiring in hiring firms (BB) that was specifically recommended to the job seeker. Lastly, column 5 focuses on hiring in hiring firms (BB) that were not directly recommended to the job seeker.

VI.2 Impact on matches: disentangling activation and targeting effects

The results of the previous section suggest that the effect of our intervention on job finding rates is not entirely driven by our targeting device. Indeed, our treated job seekers have an increased likelihood to find a job in hiring firms (BB) even when those firms were not recommended to them in the e-mails. This is not entirely surprising as our intervention was designed in a way that could very well increase the overall search effort of job seekers. In particular, we encouraged the use of the LBB platform, possibly inducing an increase in the search effort of treated job seekers directed to *any* firms presented on the LBB platform.

Our two-sided randomization design allows us to disentangle these two components of our intervention — namely, an activation and a targeting effect — by taking the analysis to a finer scale, at the level of job seeker-firm pairs. The activation effect is then defined as the increase in the likelihood that *any* match (job seeker i , hiring firm j) occurs when job seeker i is in the treated group, in the absence of any recommendation for the pair (i, j) . It captures the overall (and non-targeted) increase in search effort among treated job seekers. 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 activation effect is defined as:

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

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 her. Formally:


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

Both quantities can be identified and estimated through re-weighting strategies, and by carefully defining the population and estimation sample. Indeed, whether a given pair (i, j) was recommended or not is not orthogonal to the potential outcomes $Y_i^j(Z_i = t, R_i^j = t)$. Indeed, our recommender system was designed such that (on average) it would give higher probabilities of recommendation to pairs with a higher matching likelihood. Hence comparing recommended pairs with non-recommended pairs would not identify any causal effect. However, since we know those recommendation probabilities, we can re-weight non-recommended pairs so that their outcome distribution identifies the one of recommended pairs *had they been non-recommended*. This allows for the identification (and estimation, by the analogy principle) of the targeting effect. Similarly, the activation effect can be identified by comparing non-recommended pairs involving treated job seekers with non-recommended pairs involving control job seekers. We simply need to re-weight observations appropriately so that non-recommended pairs involving control job seekers match the distribution of non-recommended pairs involving treated ones, had the latter been control individuals.

We present the results of this exercise in Figure 3. As expected, we observe a larger activation effect among female job seekers. Meanwhile, the targeting effect appears to be stronger among

male job seekers, so that the sum of both effects yields very similar estimates for female and male individuals. Yet this analysis can still explain the large ITT among female job seekers, as for a given job seeker i , the activation effect affects the match probability with a much larger number of firms than the targeting effect. Hence overall, we would expect that female job seekers would get a larger overall effect on their job finding rate because of the larger activation effect displayed in panel (a) of Figure 3. In this respect, these results at the pair level rationalize the results presented in section VI.1.

Figure 3: ACTIVATION AND TARGETING EFFECTS, BY GENDER



Notes: This panel presents estimates of the activation and targeting effects of our intervention at the dyad level (i.e., a pair of job seeker i and firm j) on the probability that the match is actually realized — by gender. The activation effect is defined as the effect on the probability that any given match occurs when the job seeker involved in the match is treated. The targeting effect is the effect of recommending i to j on the probability that the match (i, j) occurs. The green bars reports the estimates for females, the orange bars reports the estimates for males, and the purple bars report the difference between estimates for males vs. females. The grey dashed line reports the baseline probability that job seeker i is recruited in firm j . 95% confidence intervals are reported as black error bars.

VII Social returns to the reallocation intervention: direct and indirect effects

The results of the previous section show the potential of the LBB platform in terms of directing job seekers' search toward specific firms. Going beyond this proof of concept, we ask in this

section whether this targeting effect can be used to influence the occupational scope of job search and, ultimately, to reduce congestion frictions and mismatch unemployment. As argued above, we believe this to be the key question from a policy perspective: can we in practice use recommender systems to reduce occupational mismatch, by redirecting job seekers' search toward tighter occupations? In this section, relying on Hu et al. (2022), we estimate the average direct and indirect (i.e., spillover) effects of redirecting job seekers from their occupations to neighboring ones. Ultimately, as described in section II and in Hu et al. (2022), these parameters identify a sufficient statistic for whether or not our reallocation intervention should be pushed further or not.

By reallocation intervention, we mean that our recommender algorithm generated recommendations to search in neighboring occupations when it seemed appropriate for social welfare — see section IV. For instance, job seekers in a given market m had higher chances of being recommended to apply to firms in market m' (in the same CZ, but in a neighboring occupation) if the tightness in market m was lower than in m' — see Table 5. Such suggestions of labor reallocation from slack to tight markets are made in the hope of reducing congestion frictions in slack markets while easing the matching process in tight markets. Yet at the end of the day, whether or not this led to a net social welfare gain — as measured by the overall job finding rate in the population of job seekers — remains an empirical question. It depends on both (i) the ability of the system to effectively redirect search effort, and (ii) the effect of such reallocation on the matching process in slack and tight markets.

As noted in section II and in Hu et al. (2022) it turns out that under our randomization design, one can identify key parameters to answer this question. As above, let W_i denote the indicator for whether we recommended to job seeker i any firm that was hiring in a neighboring market of i 's one. Further define for any of the $i = 1, \dots, n$ job seekers their potential outcome $Y_i(\mathbf{W}) \in \mathbb{R}$. This potential outcome is a function of the whole treatment vector $\mathbf{W} \in \{0, 1\}^n$, that gives the treatment status of all n job seekers. This underlines the fact that at this stage, we allow for any pattern of interference across job seekers. The first parameter of interest is often called the average direct effect (ADE), and is defined as follows:

$$\text{ADE} \equiv \frac{1}{n} \sum_{i=1}^n E[Y_i(W_i = 1; \mathbf{W}_{-i}) - Y_i(W_i = 0; \mathbf{W}_{-i})]$$

This parameter measures the average effect of the redirection intervention W_i on the unit being intervened on — while marginalizing over the rest of the treatment assignments of other job seekers. In a setting without interference, the ADE would match the standard average treatment effect parameter. The Horvitz-Thompson estimator for the ADE parameter is given by:

$$\widehat{\text{ADE}} = \frac{1}{n} \sum_{i=1}^n \left\{ \frac{W_i Y_i}{\pi_i} - \frac{(1 - W_i) Y_i}{1 - \pi_i} \right\},$$

where Y_i indicates whether or not i has find a job, and π_i gives the probability that W_i equals 1 — in other words, the reallocation probability for job seeker i . In our experiment, this probability is

homogeneous for all job seekers belonging to the same market, and heterogeneous across markets — slack markets having higher π 's than tight ones.

The second and less usual parameter of interest is the average indirect effect (AIE), formally defined as:

$$\text{AIE} \equiv \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} E \{Y_j (W_i = 1; \mathbf{W}_{-i}) - Y_j (W_i = 0; \mathbf{W}_{-i})\},$$

This parameter measures the average effect of W_i on all units but the one being intervened on, again marginalizing over the rest of the treatment process $\mathbf{W}_{-i}$. In other words, it corresponds to the average of the effects of job seekers' treatments on all the other job seekers. As such, it quantifies the amount of spillover effects.¹⁸ The Horvitz-Thompson estimator for the AIE parameter is given by:

$$\begin{aligned} \widehat{\text{AIE}} &= \frac{1}{n} \sum_{i=1}^n \sum_{\{j \neq i: E_{ij}=1\}} \left\{ \frac{W_i Y_j}{\pi_i} - \frac{(1 - W_i) Y_j}{1 - \pi_i} \right\} \\ &= \frac{1}{n} \sum_{i=1}^n \left\{ \left(\frac{W_i}{\pi_i} - \frac{(1 - W_i)}{1 - \pi_i} \right) \sum_{\{j \neq i: E_{ij}=1\}} Y_j \right\} \end{aligned}$$

where E_{ij} indicates whether job seeker i and j belong to the same interference space.¹⁹ The key to our design's ability to identify and ultimately yield an unbiased estimator of τ_{AIE} is the independence of the treatment status draws across job seekers — see Hu et al. (2022) appendix for further details. As demonstrated in Hu et al. (2022) appendix, the Horvitz-Thompson estimator for AIE presented in the above display is unbiased in such Bernoulli-randomized experiments.

With this in mind, we can define different variants of the AIE parameter depending on (i) the spillover effects we are interested in and (ii) the extent of the interference space that we conjecture to be relevant. We study three AIE parameters here. The first, denoted $\text{AIE}(0)$, is the average effect of reallocating i 's search effort ($W_i = 1$) on the employment outcomes of all job seekers belonging to i 's original market. Formally:

$$\begin{aligned} \text{AIE}(0) &\equiv \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i: j \in m(i)} E \{Y_j (W_i = 1; \mathbf{W}_{-i}) - Y_j (W_i = 0; \mathbf{W}_{-i})\} \\ \widehat{\text{AIE}}(0) &= \frac{1}{n} \sum_{i=1}^n \left\{ \left(\frac{W_i}{\pi_{m(i)}} - \frac{(1 - W_i)}{1 - \pi_{m(i)}} \right) \sum_{\{j \neq i: j \in m(i)\}} Y_j \right\} \end{aligned}$$

where $m(i)$ denotes job seeker i 's market. This is a key parameter as it captures the extent to which we reduce congestion frictions by redirecting search effort out of (slack) markets. On the flip side, this reallocated search effort tends to add congestion in relatively tighter markets

¹⁸In the absence of any spillover effects, we have by construction $\text{AIE} = 0$.

¹⁹In other words, if $E_{ij} = 0$, it means that we can rule out *ex ante* that j 's treatment status can affect i 's outcome in any way.

toward which it is redirected. To measure this congestion effect, we define two additional AIE parameters, AIE(1) and AIE(2), given by:

$$\forall d \in \{1, 2\}, \text{AIE}(d) \equiv \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i: j \in M(m(i)+d)} E \{Y_j(W_i = 1; \mathbf{W}_{-i}) - Y_j(W_i = 0; \mathbf{W}_{-i})\}$$

$$\widehat{\text{AIE}(d)} = \frac{1}{n} \sum_{i=1}^n \left\{ \left(\frac{W_i}{\pi_{m(i)}} - \frac{(1 - W_i)}{1 - \pi_{m(i)}} \right) \sum_{\{j \neq i: j \in M(m(i)+d)\}} Y_j \right\}$$

where $M(m(i) + d)$ is the set of markets at d steps of market $m(i)$ in the occupational graph we use. Therefore, these quantities capture the spillover effects of reallocating i 's search effort ($W_i = 1$) on the employment outcomes of all job seekers belonging to i 's neighboring markets.²⁰ Since these are the markets in which i 's search effort is redirected, we would also expect some interference to occur here. Yet in this case, we likely create some additional congestion instead of reducing them, hence we would expect some negative effect on average.

Table 7 reports estimates of all four parameters. Panel *A* presents the estimated effect on hiring outcome in any firm, while panel *B* reports the estimated effect on hiring in hiring firms (BB) specifically. The results are quite consistent across both cases. Firstly, we estimate that the ADE is very close to 0. This is interesting, as one could have feared we would deteriorate the labor market prospects of individuals encouraged to reallocate their search effort to neighboring markets. We interpret this 0 effect as implying that, from the job seeker's perspective, the cost of moving to a nearby occupation is offset by the fact that hiring prospects are better in this occupation (tighter local market). Meanwhile, the AIE(0) parameter is estimated to be strictly positive — and statistically different from 0 at the 90% level. We interpret this result as evidence that our efforts to reallocate labor out of slack markets did reduce congestion frictions to some extent in these markets. At the same time, we expect increased competition for jobs in markets toward which we redirected job seekers. Hence we need to factor in the estimates for AIE(1) and AIE(2) in our analysis to determine whether or not our reallocation intervention creates net gains in social welfare. As expected, the estimates for AIE(1) turn out to be negative — or close to 0 in the case of panel *B*. Yet their magnitude appears to be smaller than the positive effect AIE(0). Given the noise with which we estimate AIE(1) — this variance issue is even more important for AIE(2) — it is difficult to draw definitive conclusions about the overall effect of the intervention. Yet point estimates suggest that the decongestion effect in slack markets dominate the counterbalancing effect of the additional competition created in tight markets, yielding net social benefits overall. These are encouraging results for such a recommender algorithm: it suggests that pushing (at least marginally) further its reallocation component might be beneficial.

²⁰We restrict our attention to markets at 1 or 2 steps of i 's market in the occupational graph as most of the recommendations outside of i 's initial market were made toward these very close markets. As an indication, 48.9% of recommendations made in neighboring occupations were made toward markets at one step in the occupational graph.

Indeed, Theorem 1 of Hu et al. (2022) states that, in a Bernoulli trial, a positive sum of ADE and AIE implies that increasing the share of job seekers being redirected increases aggregate welfare.²¹

²¹In theory, it is not obvious up to which occupational distance one should consider that interference effects occur. Despite the fact that most recommendations were made at an occupational distance of 1, it could very well be that redirecting toward these close markets has spillover effects on the labor market prospects of job seekers that are connected to these destination markets. Hence from one market to the next, redirecting job seeker i can affect most labor markets in a given CZ in theory. Yet in practice, it is likely that such spillover effects at high occupational distance are of second order, and there is no hope to estimate those accurately, hence our choice to focus on a relatively small but relevant perimeter for the interference space. We can already observe that AIE(2) is estimated with a lot of noise in Table 7, bringing little information about the spillover effects at such occupational distance.

Table 7: ADE AND AIE OF BROADENING JOB SEARCH ON JOB FINDING RATES, BY OCCUPATIONAL DISTANCE

	ADE	AIE(0)	AIE(1)	AIE(2)
A. Effects on hiring in all firms				
Estimates	-0.005 (0.008) [0.516]	1.721 (0.991) [0.082]	-0.319 (3.69) [0.931]	-9.391 (8.087) [0.246]
Observations	441,071	441,071	430,430	423,301
Nb. local markets (CZ x Occ.)	19,511	19,511	18,152	18,383
B. Effects on hiring in hiring firms (BB)				
Estimate	-0.002 (0.004) [0.543]	0.206 (0.135) [0.129]	0.039 (0.441) [0.929]	-1.043 (0.840) [0.214]
Observations	441,071	441,071	430,430	423,301
Nb. local markets (CZ x Occ.)	19,511	19,511	18,152	18,383

Notes: Clustered standard errors at the local market level in parenthesis. P-values for the null hypothesis of a zero coefficient reported in brackets. Panel *A* presents the effects on the job finding rates in any firm, while panel *B* focuses on hiring in hiring firms only. The first column reports the average direct effect (ADE) of making at least one recommendation towards a neighboring occupation to a job seeker on its own job finding rate. The second column reports the average indirect effect (AIE) at an occupational distance of 0, i.e., the effect of making at least one recommendation towards a neighboring occupation to job seeker i on the *sum* of the job finding rates of job seekers belonging to the same local market as i . The third column report the effect on the sum of job finding rates of job seekers in job seeker i 's neighboring markets (AIE(1)) — i.e., at an occupational distance of 1 from i . The fourth column reports the same coefficient for job seekers at an occupational distance of 2 from job seeker i .

VIII Conclusion

Building upon an existing job search platform operated by the French PES, we show that recommender systems have the potential to increase aggregate employment by redirecting job seekers toward tighter occupations. This reduction in occupational mismatch comes on top of a more standard activation effect, by which recommending job seekers to make unsolicited applications to firms increases their overall search effort.

Our study uses an encouragement design, e-mailing treated job seekers. 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, a decisive feature when it comes to assessing indirect effects that are typically hard to estimate. 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. In that respect, how to make redirection suggestions salient at scale, in the typical search environment faced by job seekers — for instance by integrating to widely used platforms such tools as those developed in a more controlled environment by Belot et al. (2018b) — remains an important avenue for future research.

References

- Alfonso Naya, V., Bied, G., Caillou, P., Crepon, B., Gaillac, C., Perennes, E., and Sebag, M. (2021). Designing labor market recommender systems: the importance of job seeker preferences and competition. *Working Paper*.
- Autor, D. H. (2001). Wiring the Labor Market. *Journal of Economic Perspectives*, 15(1):25–40.
- Belot, M., Kircher, P., and Muller, P. (2018a). How wage announcements affect job search – a field experiment.
- Belot, M., Kircher, P., and Muller, P. (2018b). Providing Advice to Job Seekers at Low Cost : An Experimental Study on On-Line Advice. *Review of Economic Studies*.
- Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., and Zamora, P. (2013). Do labor market policies have displacement effects? Evidence from a clustered randomized experiment. *Quarterly Journal of Economics*, pages 531–580.
- Horton, J. J. (2017). The Effects of Algorithmic Labor Market Recommendations : Evidence from a Field Experiment. *Journal of Labor Economics*, 35(2).
- Hu, Y., Li, S., and Wager, S. (2022). Average direct and indirect causal effects under interference. *Biometrika*. asac008.
- Kircher, P. (2020). Search design and online job search—new avenues for applied and experimental research. *Labour Economics*, page 101820.
- Kuhn, P. 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.
- Kuhn, P. J. and Mansour, H. (2013). Is internet job search still ineffective? *The Economic Journal*, 124:1213–1233.
- Kuhn, P. J. and Skuterud, M. (2004). Internet Job Search and Unemployment Durations. *American Economic Review*, 94(1):212–232.
- Marinescu, I. and Rathelot, R. (2018). Mismatch unemployment and the geography of job search. *American Economic Journal: Macroeconomics*, 10(3):42–70.
- Miguel, E. and Kremer, M. (2004). Worms: Identifying impacts on education and health in the presence of treatment externalities. *Econometrica*, 72(1):159–217.

A Appendix

A.1 Context

Figure A1: LBB's HOME PAGE

La bonne boîte alternance

CONNEXION pôle emploi

Trouvez ici les entreprises qui recrutent régulièrement, et contactez-les !*

Métier recherché : (boucher, cariste, secrétaire, ...)

Autour de : (Paris, Bd Voltaire, 33000...)

Vous recherchez un contrat en alternance ? Les entreprises susceptibles de vous recruter sont sur [La Bonne Alternance](#).

*Grâce à un algorithme exclusif de Pôle emploi détectant les entreprises qui vont probablement embaucher ces 6 prochains mois.

Conseils
Code source ouvert
A.G.P.D.
Accessibilité

F.A.Q.
C.G.U.
Accès recruteurs

API
Espace Presse
Contact

UNION EUROPÉENNE

Ce dispositif est cofinancé par le Fonds Social Européen dans le cadre du Programme opérationnel national "emploi et inclusion" 2014-2020

Figure A2: LBB's RESEARCH RESULTS PAGE

Enseignement supérieur Paris 75001

181 entreprises sont susceptibles de recruter en Enseignement supérieur autour de Paris

Masquer la carte

Trier

Tri optimisé Distance

Affinez votre recherche

Secteur d'activité

Tous les secteurs

Taille de l'entreprise

Toutes tailles

Moins de 50 salariés

Plus de 50 salariés

Distance

5 km

10 km

30 km

50 km

100 km

+ de 100 km

UNIVERSITE PARIS 1 PANTHEON-SORBONNE - PARIS-05

Enseignement supérieur

500 à 999 salariés

2,2 km de votre lieu de recherche

Potentiel d'embauche

★★★★☆ (4)

Plus d'infos Enregistrer dans MEMO Postuler

UNIVERSITE PARIS DIDEROT - PARIS 7 - PARIS-13

Enseignement supérieur

250 à 499 salariés

Potenti

Donner votre avis

Figure A3: LBB's FIRM CONTACT INFORMATION PAGE

UNIVERSITEPARIS1PANTHEON-SORBONNE - PARIS-05

Enseignement supérieur
500 à 999 salariés
2.2 km de votre lieu de recherche

Potentiel d'embauche
★★★★☆

[Plus d'infos](#) [Enregistrer dans MEMO](#) [Postuler](#)

Raison sociale
UNIVERSITE PARIS 1 PANTHEON
SORBONNE

Contact
racbiatss@univ-paris1.fr
0144077918

Mode de contact à privilégier
Envoyer un CV et une lettre de motivation

Enseigne
UNIVERSITEPARIS1PANTHEON-SORBONNE

C'est mon entreprise !
[Modifier ces informations](#)

Informations supplémentaires
[Google](#)
[Kompass](#)
SIRET : 19751717000019

Adresse
Service des ressources humaines
12 PLACE DU PANTHEON
75005 PARIS-05

[Télécharger la fiche en PDF](#) [Donner votre avis](#)

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 Occupational distance and observed transitions

Figure A5: MEAN OCCUPATION DISTANCE VS OBSERVED RANK IN OCCUPATIONAL TRANSITIONS



Note: This graph 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.3 Commuting zones and local labor markets

A.3.1 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.²² 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 A1 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.

Table A1: COMMUTING ZONES' STATISTICS

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

Standard errors in parenthesis.

²²We miss data for one commuting zone which regroups Saint-Martin and Saint-Barthélemy.


A.3.2 Local Labor Markets

Upon registering 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 services²³). We define a local labor market as the intersection between commuting zones and occupations. In France there are 404 CZ and 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.

²³ROME stands for "Répertoire opérationnel des métiers": Operational directory of occupations.

A.4 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.

$$\text{LOG}(\text{REALIZED HIRINGS}) = 1.33(0.0053) + 0.89(0.0039) \times \text{LOG}(\text{PREDICTED HIRINGS}), R^2 = 0.37$$

A.5 Ex-post treatment

Figure A7: NUMBER OF DISTINCT CLICKS BY TREATED ESTABLISHMENT



Note: Distribution of the number of distinct clicks (one per job seeker) per establishment. The displayed distribution is cut above the 99th percentile. The average number of distinct clicks per establishment is 9.1

Table A2: OVERALL NUMBER OF CLICKS FOR ESTABLISHMENTS IN COMMUTING ZONES WHERE 60% OF FIRMS WERE TREATED

	(1)	(2)	(3)
	Pre intervention	During intervention	Post intervention
ITT	0.0124 (0.0908)	1.539 (0.0761)	0.0211 (0.0547)
Constant	3.912 (0.143)	1.590 (0.0635)	1.864 (0.0751)
N	47305	47305	47305
Mean	3.920	2.516	1.877
Adjusted R^2	-0.0000208	0.0100	-0.0000182

Note: ITT of the overall number of clicks for establishments in commuting zones with a 60% treatment rate during (1) the pre-intervention period, (2) while the intervention is going on and (3) in the month following the end of our intervention. Regressions are weighted by inverse treatment status probability. Standard errors are clustered at the labor market (Sector*CZ) level.

Table A3: OVERALL NUMBER OF CLICKS FOR ESTABLISHMENTS IN COMMUTING ZONES WHERE 20% OF FIRMS WERE TREATED


	(1)	(2)	(3)
	Pre intervention	During intervention	Post intervention
ITT	0.0221 (0.114)	2.044 (0.114)	0.0820 (0.0601)
Constant	3.311 (0.0849)	1.539 (0.0422)	1.548 (0.0399)
N	51061	51061	51061
Mean	3.315	1.951	1.565
Adjusted R^2	-0.0000185	0.0206	0.0000337

Standard errors in parentheses

Note: ITT of the overall number of clicks for establishments in commuting zones with a 20% treatment rate during (1) the pre-intervention period, (2) while the intervention is going on and (3) in the month following the end of our intervention. Regressions are weighted by inverse treatment status probability. Standard errors are clustered at the labor market (Sector*CZ) level.

A.6 Additional results on private returns

Figure A8: JOB-FINDING RATE ITT ESTIMATES




Note: This graph presents the ITT estimates for job finding at different time horizons. Sample restricted to job seekers who were still unemployed as of 19/11/2019. Standard errors are clustered at the labor market (Occ.*CZ) level and associated 95% confidence intervals are displayed.

Further decomposing women's response into access to indefinite as opposed to definite duration employment (Figure A9), we find that the positive effect of our intervention is driven by a rise in treated women's return to definite duration employment.²⁴

²⁴A further decomposition between "long term" (i.e. more than six months) definite duration contracts and short term (i.e. less than six months) definite duration contracts shows that this effect is driven by short term definite duration contracts.

Figure A9: JOB-FINDING ITT ESTIMATES BY CONTRACT TYPE FOR FEMALE MALES



Note: ITT estimates for job finding of (a) indefinite duration and (b) finite duration contracts at different time horizons. Sample restricted to female job seekers who were still unemployed as of 19/11/2019. Standard errors are clustered at the labor market (Occ.*CZ) level and associated 95% confidence intervals are displayed.

Potential mechanisms underlying gender differences

Differences in observable characteristics

Women's and men's responses to tailored job-search advice appear to be strikingly different. Could this difference be driven unbalances in the gender distribution across observables and labor markets? In other words, are women reacting more to our treatment because they differ in some observable way from men or because they work in occupations that tend to respond more strongly to the provision of tailored job-search advice. To check this, we interact our intention-to-treat status with a male/female dummy and control for the interaction of our treatment with a set of observables, including a full set of labor market fixed effects. We present the results of these robustness checks for definite duration hirings in Table A4. The different response of men and women stays remarkably robust for all the interacted controls and interacted labor market fixed effects we include, indicating that the gender differences in the response to our provision of tailored job search recommendations do not appear to be driven either by individual level observables being correlated to gender differences or by labor market differences.

Table A4: ROBUSTNESS CHECK: DO DIFFERENCES IN OBSERVABLES EXPLAIN TREATMENT EFFECT HETEROGENEITY ACROSS GENDER

	(1)	(2)	(3)
Male # ITT	-0.0420 (0.135)	-0.0367 (0.135)	-0.221 (0.149)
Female # ITT	0.287 (0.108)	0.309 (0.110)	0.257 (0.130)
Controls	No	Yes	Yes
Labor Market FE	No	No	Yes
Observations	800297	800237	793103
Mean	0.154	0.154	0.154
Adjusted R2	0.00201	0.0203	0.109

Note: This table displays the results of a regression of finite duration job-finding on the interactions of our treatment with a dummy for males and a dummy for females. Column (1) does not add any control, column (2) controls for the direct and interacted effects of the centered value of age, a diploma dummy, experience and unemployment spell duration. Finally column (3) adds the direct and interacted effect of centered labor market (Occ.*CZ) fixed effects calculated through a first stage regression. Sample restricted to job seekers who were still unemployed as of 19/11/2019. Standard errors in parentheses are clustered at the labor market (Occ.*CZ) level. Coefficients and standard errors in percentage points.

Differences in take-up

To further investigate the gender differences in job seekers' responses to our intervention we try to follow gender differences along the causal chain that eventually links our intervention to the hiring of a job seeker. This causal chain starts with opening of emails, then goes on with clicking on links, applying to firms, being called for an interview, receiving an offer, accepting it. We start from the beginning by first looking at gender differences in initial take-up measures. To do so we regress our main take-up measures, opening at least one email and clicking on at least one link, on a male/female dummy. Table A5 shows that men are 6% less likely to open the emails we sent them. This big difference in take-up passes through to subsequent clicks and remains large when we include detailed individual level controls as well as labor market fixed effects. The fact that women are 25% more likely than men to click on the recommendation link we sent them cannot, however, fully account for the gender differential we see on final outcomes. The initial variation in take-up must hence be complemented by other differences involving latter stages of the hiring process. Unfortunately we were not able to track applications and interviews of all treated and control job seekers. One possibility could for instance be that men and women

react differently to suggestions to widen the occupational breadth of their job-search effort — we investigate this possibility in the following subsection exploiting our web survey.

Table A5: GENDER DIFFERENCES IN TAKE-UP (in percentage points)

	Opened email			Clicked on link		
	(1)	(2)	(3)	(4)	(5)	(6)
Male	-6.733 (0.294)	-6.645 (0.250)	-3.982 (0.189)	-5.957 (0.258)	-5.796 (0.253)	-3.458 (0.174)
Controls		Yes	Yes		Yes	Yes
Fixed effects			Yes			Yes
N	533557	533557	525702	533557	533557	525702
Mean	0.639	0.639	0.639	0.245	0.245	0.245

Standard errors in parentheses

Note: Regression of (1,2,3) opening at least one email and (4,5,6) clicking on at least one link on male female dummy. We add individual level controls in columns (3,4,5,6) as well as labor market fixed effects in columns (3,6). Sample restricted to treated job seekers who were still unemployed as of 19/11/2019. Standard errors are clustered at the labor market (Occ.*CZ) level. Coefficients and standard errors in percentage points.

A.7 Choice of the screening technology:

More specifically we choose to parametrize our screening function $q^{f,j}$ as:

$$q^{f,j}(\theta^{f,j}) = \frac{1}{[1 + (\frac{\theta^{f,j}}{\Gamma m_f \bar{\theta}_j})^\gamma]^{1/\gamma}}$$

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

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

And where $\bar{\theta}_j$ denotes the local slackness ratio in occupation j . This local slackness ratio is defined as the ratio of possible recommendations present in the vicinity of occupation j to the total number of hirings predicted in occupation j . Formally:

$$\bar{\theta}_j = \frac{\sum_w \rho_w^{d_{i(w),j}} T(w)}{\sum_f V^{f,j}}$$

For $\gamma > 1$ this function is monotonous in $\theta^{f,j} = W^{f,j}/V^{f,j} > 0$ and verifies:

$$q^{f,j}(0) = 1$$

$$q^{f,j}(+\infty) = 0$$

What's more $q^{f,j}$ has an inflection point at $m_f\theta_j$ so that according to the value of m_f , firm's f congestion effect will start to quick in either before ($m_f = m_f^L < 1$) or after ($m_f = m_f^H > 1$) the number of recommendations sent to (f, j) relative to its predicted hirings (i.e $W^{f,j}/V^{f,j}$) reaches the local slackness ratio θ_j .

In practice we select:

$$m^L = 0.5$$

$$m^H = 1.5$$

$$\gamma = 3$$