

JOB SEARCH IN THE 21ST CENTURY

Philipp Kircher

REPRINT | 3275

CORE

Voie du Roman Pays 34, L1.03.01

B-1348 Louvain-la-Neuve

Tel (32 10) 47 43 04

Email: lidam-library@uclouvain.be

<https://uclouvain.be/en/research-institutes/lidam/core/core-reprints.html>

JOB SEARCH IN THE 21ST CENTURY

Philipp Kircher

Cornell University, USA, Université
Catholique de Louvain, Belgium, and
University of Edinburgh, UK

Abstract

Formal job search has undergone massive transformations in the last two decades with the move to online platforms. This makes more stages of the job search process observable, and allows researchers to intervene in this market and improve the design of the job search process and the information available in it. This paper showcases studies that illustrate the potential of analyzing online job search data and of intervening in the online job search process. It presents the current insights and some open questions and highlights conditions under which some of the recent interventions are likely to improve market outcomes overall, rather than improving only the outcomes for the treated individuals. (JEL: J64, J62, C93)

1. Introduction

Formal job search has changed massively over the last two decades of this millennium. At the end of the last millennium, less than a quarter of unemployed Americans searched for a new job online. Arguably, this was due to a lack of high-speed Internet access, since the majority of unemployed Americans who had high-speed Internet

The editor in charge of this paper was Andrea Weber.

Acknowledgments: The content of this article is based on the Schumpeter lecture presented at the European Economic Association Annual Meeting in Milan in 2022. I am grateful to the large number of co-authors who have contributed to the research agenda on interventions on job search platforms. This includes special thanks to Michèle Belot and Paul Muller with whom I stumbled into this research agenda a decade ago and who have inspired his thinking and the direction of research ever since. I also thank the European Research Council (grant 284119) and the UK Economic and Social Research Council (grant ES/L009633/1) for financial contributions to this research endeavor, and the Université Catholique de Louvain and the University of Edinburgh for their continuing support. Invaluable research assistance has been provided by Gabriele Macchi for the proofs in this paper, and by Elmer Zongyang and Nathan Mislav in the preparation for publication. The author has no conflicts of interests, but simply notes that the co-authors of the paper Horton, Johari and Kircher (2021) reviewed in this article were involved in various roles (consultant/adviser/employees) at the company where the study was conducted.

E-mail: pk532@cornell.edu

Journal of the European Economic Association 2022 20(6):2317–2352

<https://doi.org/10.1093/jeaa/jvac057>

© The Author(s) 2022. Published by Oxford University Press on behalf of European Economic Association.

All rights reserved. For permissions, please e-mail: journals.permissions@oup.com




FIGURE 1. Internet access and job search among US workers (%). Source: Kuhn (2014, Figure 1).

access at that point already used it for formal job search. One decade later, more than three-quarters of unemployed Americans use the Internet for job search, with even higher numbers for those who have access to high-speed Internet (see Figure 1 based on Kuhn (2014)).

This change in the nature of job search offers two major opportunities: First, it allows researchers to study the job search process in unprecedented detail. For a long time, the matching between workers and firms has been treated as a “black box” within the matching function, because insights from existing survey data were too limited to discipline the exercise. The advent of matched employer–employee data revolutionized our understanding of the labor market, because extensive information on the existing matches became available. Still, these datasets only contain information on the matches that actually formed, but not on the process by which they came about. New datasets on online vacancy postings and online job searches allow a detailed picture of what gets announced, which jobs workers consider and apply for, and possibly which workers are selected for interviews and eventual hires. This is likely to generate a new revolution in the way we study job search and the formation of new matches in the labor markets.

Second, and maybe even more exciting, the rise of online job search might fundamentally alter the way workers and firms match. At the turn of the millennium, early successful platforms like Craigslist mostly replicated existing classified sections in newspapers: One could scroll through jobs in one’s geographic neighborhood, akin to searching in a local newspaper, only at lower cost.¹ Two decades later, such platforms offer a much more sophisticated search interface that allows job seekers to specify keywords, location, and radius to explore the wealth of existing job listings (see Figure 2).²

1. Craigslist already had a keyword search on its central home screen, but it was clearly not the main way of searching, and it made it difficult to arrange keywords so that it returns relevant jobs in the local geographic area.

2. Craigslist is an interesting case because it was a market leader, and its rollout across the United States has been used as an instrument to study the effects of the move online on various domains. Kroft and Pope




FIGURE 2. Search interface of Craigslist, 2000 versus. 2022. Source: The view for Craigslist in the year 2000 was obtained from repository <https://web.archive.org/web/20001109014500/http://www.craigslist.org/sfo/acc/>. The view for Craigslist in the year 2022 when taken by screenshot from their website. Both were taken in August 2022.

Still, the recent improvements seem to be far from exhausting the potential that such platforms have.³ Because information provision and customization are so powerful and cheap online, one can in principle tailor the job search experience at the individual level. Therefore, economists can contribute their knowledge of labor markets to aid in this process, opening the avenue to become market designers instead of just observers. Both academics as well as private job search platforms are trying—sometimes in collaboration—to improve the job search process through more and better information, easier interaction, and A/B testing, which allows for rigorous evaluation.⁴ Redesigning the interactions and information in these markets give rise to many novel questions, and this article provides some recent approaches on the academic side and some general considerations for this process.

(2014) document little effect on labor markets, which might be partially due to the fact that the platform did not initially adjust its delivery method for classifieds to the new online medium—even though the long-term impact after such adjustments have been made might be large.

3. Craigslist is an example of a platform that provides particularly little advice to job seekers. There are other platforms that use algorithms to recommend jobs, but finding an optimal system poses novel challenges in the job market setting, which are discussed in Section 4.

4. A/B testing is a shorthand for a randomized controlled trial of platform features with subsets of the users.

This article briefly highlights some existing studies that use observable information on labor markets to shed new light on the job search and matching process (Section 2). The remaining chapters deal with interventions in online job search markets, starting with some general examples (Section 3) but then focusing on the provision of advice: Section 4 discusses some of the considerations when choosing the type of advice to be offered, while Section 5 goes deeper into some examples of studies that implemented advice. It also discusses advice provided outside of job search platforms and the costs and benefits thereof. Section 6 provides a simple illustrative model that highlights the settings in which changes in search strategies that benefit treated individuals are likely to generate positive aggregate effects. This provides guidance about the type of settings in which effects of large-scale rollouts are particularly likely to yield positive outcomes. Section 7 concludes. While some sections review existing debates, novel considerations are more concentrated in Sections 4, 5.2 and 5.3, and 6. Overall, the paper aims not only to showcase the current stance of the literature, but also to discuss open questions that remain.

2. Using *observables* to Study the Black Box of the Matching Function

Studies that use data from online job search sites to open the black box of the labor market matching function have grown rapidly in recent years. This section reviews just a few of them to showcase the type of insights that this new data can generate. For more details, see, for example, the discussion of online search in Kircher (2020).

A nice set of recent studies illustrate how observing the various stages in the search for workers (by firms) and for employers (by workers) can be used to shed light on important questions related to gender: Kuhn and Shen (2013) study the request for a particular gender in Chinese job postings. Because the postings are online, they can relatively easily be scraped. Then, text analysis can be applied to gain insights into the demands of employers. The paper shows that 10% of jobs in their sample are gendered, and roughly half of these ask for men and the other half for women. While this might be encouraging, job advertisements that also feature the words “young” and “beautiful” are nearly exclusively looking for female employees, while advertisements that ask for experience or seniority nearly exclusively ask for males. Kuhn, Shen, and Zhang (2020) built on this to explore the next step in the interaction: the employers response rate to workers who applied to gender-mismatched job postings. While employer responses are often not recorded on the platform, progress in the ease of interacting directly on the platform lowers the costs for firms and renders it likely that such interactions can soon be completely tracked on such sites. Finally, Kuhn and Shen (2021) revisit this after gender was prohibited as a criterion in job postings in China.

This shows that not only the job postings and the worker responses can be studied, but also subsequent call-back and possibly hiring. Still, the vast majority of studies have exclusively concentrated on the worker side. These studies ask how workers search, and differ according to the dimension along which this is explored. They also differ in the combination of datasets, ranging from exclusive consideration of online

interactions, to studies that combine this with survey methods, field experiments, or administrative data. The following gives some examples of the possibilities.

Job Search by Discount Factor. Belot, Kircher, and Muller (2022b) develop a methodology to elicit short- and long-run discount factors of job seekers. They combine a field experiment that elicits these with observations of online search behavior, which they correlate and compare to existing model predictions.

Job Search over Time. Marinescu and Skandalis (2021) link online job search effort to administrative data on unemployment benefits. They document a gradual increase in search effort up to benefit exhaustion and a gradual decrease thereafter, a pattern that they then use to distinguish between different models of job search. Faberman and Kudlyak (2019) also investigate the change in search effort over time, albeit without the ability to link it to administrative data.

Job Search over Space. Marinescu and Rathelot (2018) document the spatial radius of applications on a large US online platform, and show that the vast majority of applications are very local, though some span long distances. Interestingly, the spatial distribution is virtually identical for those with different levels of education. This type of study provides insight on how to think about separate labor markets within a larger geographic entity such as the United States.

Job Search over Salaries. Marinescu and Wolthoff (2020) document that jobs which pay more get fewer applications even after controlling for occupations at a granular level and even after adding firm fixed effects. This negative relationship reverses to a positive elasticity of roughly 7% more applicants per 10% higher wage, when controlling for job titles and firm fixed effects. This suggests that occupations are not sufficient controls to ensure that jobs are comparable. Occupations miss notions of seniority that are captured by job titles, such as “senior” administrator. Banfi and Villena-Roldan (2019) replicate this in a developing country. They additionally show that wages predict application behavior after controlling for job titles, even when the wages are hidden from prospective applicants. Therefore, some parts of the job description beyond the job title still seems to vary between vacancies, and is not controlled for by job titles. Obviously, in many labor market models, the role of wages in attracting applicants is very relevant and clear empirical estimates for this are relevant.⁵

These studies show that one can get relevant insights into how individuals and firms search, which can be used to discipline underlying models and to generate new hypotheses for these markets. They also show the usual shortcomings of observational data, which often lack the random variation that would enable clear identification.

5. Mueller et al. (2018) consider a similar question from the firm side and ask how changes in the wage impact the speed at which firms can fill vacancies by linking online job posting data to administrative matched employer–employee data.

Currently, insights on the content of job postings and the search and application behavior abound, while later stages of the recruitment process have so far received much less attention. These include the call-back from employers, the wage offers, and the final acceptance. As platforms expand their services and make it ever easier and more cost-effective to interact on the platform rather than outside of it, these later stages might soon become more observable and might lead to a new research focus in this literature.

3. *Intervening in the Online Job Search Market-Place*

Rather than acting as observers, researchers are enabled by the online marketplace to directly affect and potentially design the market, and to test the interventions in rigorous controlled trials. Here we present some background and a few examples in this literature that is still developing. As more academics engage in this area, and as more platforms do A/B testing with economists involved in the design and analysis, we expect this literature to grow rapidly, despite its current small size.

Together with the excitement associated with a direct role of researchers in the job search process come conceptual and ethical issues. On the conceptual side most research studies in this area have not yet addressed general equilibrium effects: Jobs tend to be scarce, and helping some individuals to find jobs might take these jobs away from others. The next section discusses this issue and other conceptual points in some more detail in the context of online job search advice, and Section 6 is entirely devoted to this point: It provides a theoretical discussion of the kinds of settings in which the benefits to a treatment group are likely to have positive effects even in the aggregate when spillovers on other individuals are taken into account. This might help to guide future implementation efforts.

On the ethical side, unemployed individuals are in a vulnerable position. Therefore, studies should not place a burden on job seekers, and possible negative consequences should be mitigated to the greatest extent possible. For some types of studies this is more easily achieved. For example, in most studies that rely on information provision on search platforms, the information can simply be disregarded. And if the information is taken on board, then it was designed with the aim to be helpful. In other settings, the ethical concerns are substantially more challenging, and designs that minimize such concerns are especially relevant. The following study provides one such example.

Belot, Kircher, and Muller (2022c) take up the question that concluded the previous section, and randomized the wage on fictitious vacancies. This ensures that the wage is not correlated with other parts of the job description.⁶ They embed these fictitious

6. Dal Bó et al. (2013) also randomize the wage. They do so for real jobs, which leads them to focus on a single type of job in the public sector of a developing country. They do their randomization at a later stage in the recruitment process, and highlighting the role of mission motivation. By using fake instead of real jobs, the randomization can span a much larger set of occupations, but the ethical considerations require a very careful approach.

vacancies in a large number of real vacancies, and job seekers know that a very small fraction of fictitious vacancies exists but not exactly which ones. While searching, job seekers can save vacancies to which they intend to apply later. Before they actually start applying for jobs, they are informed if any fictitious vacancies are among the ones that they have saved. These features ensure that job seekers are not deceived, and that the costs to them are kept to a minimum.

The study finds that a 10% wage increase is associated with 7%–9% more intended job applications. This is very close to the observational studies that controlled for job titles, though the context and country are very different. The vacancies they post include pairs that are essentially identical except for the wage. The study finds that job seekers who apply to the low-wage job often do not apply to the high-wage job (around 40% do not apply to the higher wage). This persists even after controlling for both jobs being displayed to them on the screen. Such evidence opens new questions about who selects into low-wage jobs even though identical higher-wage jobs are available in the market: For example, could there be gender-differences in the avoidance of more competition at higher wages that could explain this? Their design is underpowered to explore this question, though larger studies should be able to uncover possible differences across job seekers. In general, their design allows them to flip the classical audit studies where fake resumes are sent to employers, and confront job seekers to see their reactions to fake jobs that differ in one dimension only. Their design shows a way to mitigating the ethical issues involved.

Online marketplaces also have the ability to provide real-time data that classical platforms such as newspaper job-boards would have struggled with. Related to the question of gender, Gee (2019) explores the role of showing job seekers how many other individuals have already applied for a job, and considers the differential impact by gender. Since individuals can concentrate on jobs that do not yet have too many applications, application rates increase. The effect is mostly concentrated on women, whose application rate improves substantially. Relatedly, Horton (2019) offers workers a costly signal that they can use to show employers that they are ready to take on work, and the paper shows that this leads to more employer contacts, less rejected job offers, and higher market surplus.

This touches on a particularly relevant question from a market-design perspective: How a platform can collect new match-relevant data either from users or from outside sources, and use it for match-making. A particularly low-cost intervention is cheap-talk. Horton, Johari, and Kircher (2021) ask employers who post a vacancy to declare what they are interested in: an entry-level worker at low pay, an expert at high pay, or someone in the middle at middle pay. Figure 3 shows the actual choice scenario. They randomize whether this information is revealed to prospective workers or not. Even without revealing the choice, the quality (as measured by previous experience on the platform) of applicants at entry-level jobs is lower than the quality of applicants at expert jobs. This might not be surprising as firms can indicate their preferences in the job description. But this amplifies once the information is revealed. The reason might be that the information now becomes more standardized, visible, and searchable. Moreover, workers lower their wage demands when an “entry-level” declaration is

ENTRY LEVEL	\$	INTERMEDIATE	\$\$	EXPERT	\$\$\$
I am looking for freelancers with the lowest rates		I am looking for a mix of experience and value		I am willing to pay higher rates for the most experienced freelancers	

FIGURE 3. Vertical preference signaling opportunity presented to employers. Source: Horton, Johari, and Kircher (2021).

revealed rather than hidden, and increase them when an “expert” declaration is revealed rather than hidden. The authors then analyze the implications for truth-telling and the efficiency of matching.

Relying on such cheap talk is a low-cost way to generate more information for the market and facilitate matching. In principle, cheap talk could be used to inquire about many dimensions in the firm-worker interaction. Horton, Johari, and Kircher (2021)’ study is one relevant example; alternatively, one could for example ask employers about their urgency of hiring, their interest in employing workers with unconventional backgrounds, or the likely job duration, which could be displayed to the workers. Finding the most relevant dimensions, selecting those where truth-telling is plausibly a profitable strategy for firms, and keeping complexity of the questions and answers low, will likely be relevant for understanding the full benefit of this approach.

The cheap talk intervention opens the question more generally of whether there is information that is readily available to labor economists that would benefit market participants in the job search process. In the previous interventions, information was collected and made available to the other side. It was not directly labeled as advice. Since job search advice could be a key advantage of online platforms, we review this here in more detail. The following section provides a short discussion of possible guidelines for providing such advice effectively, before highlighting a few studies with particular implementations.

4. Intervening through *Advice* —What Advice?

An algorithmic recommender system can provide job seekers with suggestions of vacancies they should look at. The question is, on what information such recommendations should be based on. The labor market does not function like the market for books or podcasts. When Spotify recommends “Think Like an Economist” to those looking for “podcasts in economics” because other people interested in podcasts in economics listen to this one, they create additional demand for this podcast. This demand can easily be satisfied. This is where labor markets differ. It might not be wise to show a particular vacancy to additional job seekers when many other job seekers are interested in it, because usually only one of them can get the job. This scarcity

problem at the job level creates additional complexity. Providing vacancy-level advice generated through machine learning might also make it difficult for the job seeker to understand exactly what the advice is. Seeing many different recommendations, the job seeker might not be able to tell the common thread, which might make it hard to use the information outside of the platform. While some of these considerations have just entered the discussion of economists in the context of algorithmic job recommender systems (see, e.g., Naya et al. 2021), the question how easily humans can learn the algorithmic advice and use it in other settings is not yet part of the broader discussion.

A longer-existing alternative is to use much coarser categories of advice. A particular dimension that has been explored is advice about occupations or sectors. There are three advantages of using this approach: (1) Even though scarcity also applies at the occupational level (one does not want to send all workers to the same occupation), labor economists already have some rough understanding of the types of settings in which this is less of an issue. (2) Since there are not that many occupations, it is relatively easy for workers to learn and then to apply this knowledge more broadly when searching for jobs, for example by talking to friends and family. (3) Occupational advice can be provided relatively easily, though more sophisticated machine learning algorithms might improve this further.

For the first point, economists have long pointed out that it is inefficient (all else equal) to have more job seekers per vacancy in one market than in another. To the extent that job seekers can move from one market to the other at relatively similar productivity, this tends to improve matching. Inducing more search also tends to improve labor market matching. So does shifting job search to sectors or occupations where the worker has an individual advantage in finding a job. Finally, moving work from short- to long-term unemployed individuals is beneficial as the former can find alternatives more easily. While these insights are intuitive and widely known, the author does not know a reference that easily spells them out. In this regard, Section 6 is devoted to a stylized toy model that illustrates these forces and highlights the caveats.

For the second point, it is useful to recall that around half of the jobs that unemployed job seekers obtain come through referrals. For online advice to be particularly useful, it is crucial that individuals not only use it when they engage in formal job searches but also when they turn to their social networks. This requires job seekers not only to consider recommended vacancies on an online platform, but to actually develop a deeper understanding of the advice that is provided. For this, it might be essential that advice is grouped into large aggregates, like occupations. As an example, if an economist loses his job and is recommended to look for a job in banking, he can use this information to approach those in his social network that have already worked in banking. Advice that comes out of a recommender system might not have such a simple common thread, and therefore might be harder to understand. Still, it might be crucial to embed even the simple occupational advice into a job search engine, because showing the actual vacancies associated with “banking” might make the advice less “academic” and more real, and therefore easier to learn.

Finally, it is relatively simple to start generating occupational advice. Good advice needs at least two features: The recipient should have the skills to work in the

TABLE 1. Occupation transition matrix (fictitious example).

Occupation in 2015	Occupation in 2016		
	Waiter	Retail Sales Assistant	Bar Staff
Waiter	80%	15%	5%
Retail sales assistant	15%	80%	5%
Bar staff	5%	5%	10%

recommended occupation, and there should be jobs available there. A simple way to generate advice is to ask individuals about their previous occupations (or the occupations they are looking for) and to see where those who work in those occupations transit to. Such a transition matrix is illustrated in Table 1. It can be interpreted as follows: 80% of those who worked as waiters and waitresses in year 2015 still did so in 2016, while 15% moved to “Retail Sales Assistant”, 5% to “Bar Staff”, and essentially none became economists. Therefore, waiters and waitresses seem able to do the job of retail sales assistants, and there seem to be enough jobs out there to switch into, making it a promising recommendation.

With a large enough population, in principle the transition matrix could be obtained for subsets of the population. For example, transitions could be observed for individuals with particular labor market experience, gender, and age. It could then be used to provide advice for others that fall into the same subgroup. While such advice might be more targeted, it risks perpetuating existing patterns into the future. More sophisticated machine learning algorithms share similar concerns about the perpetuation of biases. By using a more basic transition matrix that does not condition on those finer categories, such biases can be avoided. The obvious cost is that the advice might be less clearly targeted.

Beyond the transition matrix, one can simply look for skill transferability, which sites like O*NET provide. This might recommend occupations might have low labor demand, though. One solution is to supplement this with a display of market tightness (the ratio of job seekers to vacancies in an occupation) by geographical area to indicate job availability. One can also display occupations that require little skills and have high demand, or additional information such as the risk of future automation, etc. Many of these approaches have been explored as parts of the following studies, though the transition matrix arguably provided the most visible advice therein.

A careful analysis of which pieces of advice are most useful, and to which extent this is context-dependent, remains an interesting question for future work. A deep answer to this probably has to weigh the benefits of many pieces of advice against the costs of increased information overload. Another open question is how simple methods (like those based on the transition matrix) compare to more complex machine learning methods. Even when restricted to providing only occupational advice (rather than recommending individual vacancies) machine learning could incorporate more

information than that of a simple transition matrix.⁷ As discussed two paragraphs earlier, the benefits of better advice then need to outweigh the potential risk of perpetuating existing biases.

5. Intervening through Advice —Examples

5.1. A Small-Scale Study of Occupational Advice Online

Belot, Kircher, and Muller (2019) recruited roughly 300 job seekers in Edinburgh in 2013–2014, who were followed for 12 weeks and were recruited to search on a job search website that was professionally built for that purpose. It allows searches over a vacancy feed covering more than 90% of the UK vacancy count. For the first three weeks, all participants had to use keyword search. This means that they had to come up with keywords themselves, specify a location and radius, and were shown the most relevant responses, often in the order of posting data. Such an approach is still one of the most common ways to search for jobs online, as exemplified in Figure 2 in the Introduction. Since individuals do not yet know their treatment status, one can measure how many occupations they search by computing a measure of occupational distance between jobs according to a metric on the occupational codes.

In week 4, half of the participants were assigned to a treatment group, where they were asked about the occupation they were looking for, after which they were then presented with a list of related occupations as illustrated in Figure B.1 in Appendix B, based on the transition matrix as well as on skill transferability. They could view the market tightness in geographical heat maps, and by pressing “search”, they saw the resulting vacancies for all the recommended occupations in their area. The vacancies were displayed in batches to fill the screen, and applicants had to manually move forward to see items further down the line. Those items that actually made it onto the screen are “listings”, which could be viewed by job seekers for more information, and saved to apply later. For each saved vacancy, job seekers were asked in subsequent weeks whether they had applied, obtained an interview, or got hired. They were also asked about the total number of applications and interviews elsewhere.

To see if this broadens the occupational breadth of job seekers, the paper constructs an index of the distance between any two vacancies based on differences in the occupational code. The occupational breadth is then the average occupational distance among all vacancies that are listed/applied for in a given week by a particular job seeker. For some of the results, it is useful to consider individuals who in the first three weeks (prior to the intervention) search narrowly (defined as searching with average occupational breadth below the mean) and those that search more broadly (the complement).

7. Machine learning could also directly select individual vacancies, but as noted earlier, this opens questions about the scarcity of hires for any particular job and whether jobs seekers can learn such information and use it in other job search activities outside the platform, such as talking to friends and family.

TABLE 2. Effect of intervention on the occupational breadth of listed vacancies and applications.

	Occupational Breadth	
	listings	applications
Treatment	0.13** (0.06)	0.03 (0.20)
Treatment X occupationally broad	-0.07** (0.04)	-0.43* (0.22)
X occupationally narrow	0.34*** (0.10)	-0.49* (0.29)

Source: Belot, Kircher, and Muller (2019) Tables 6 and 8. Numbers represent point estimates from a linear difference in difference regression, brackets provide standard errors, stars indicate 10% (*) and 5% (**) significance, and occupationally broad refers to the subject pool who had above average occupational breadth in the 3 weeks before the intervention, while occupationally narrow refers to the complement.

Table 2 replicates results from Tables 6 and 8 in Belot, Kircher, and Muller (2019) and shows the impact of the intervention on the treatment group compared to the control group. The treatment group lists a significantly broader set of vacancies on their screen. People also become broader naturally over time. To interpret the magnitude, the intervention increases occupational breadth by an amount that is similar to the natural increase in breadth over the twelve weeks of the study within the control group. Not surprisingly, the effects are larger when comparing narrow individuals in the treatment group with narrow individuals in the control group. For occupationally broad individuals, the effects reverse, and the treatment group becomes narrower than the control group. That might be intuitive, as those who consider nearly every job in Edinburgh might become more focused as part of the intervention. While these effects on listed vacancies are somewhat mechanical, they translate into changes in the occupational breadth of applications for the different subgroups. This shows a causal effect on their application behavior.

There is little effect on other outcomes, such as the number of listings or the number of applications. There are indications that the geographical breadth declines when occupational breadth increases, which might be because the increased set of vacancies allows people to search closer to home.

Table 3 captures the effect on job interviews and replicates Table 10 in Belot, Kircher, and Muller (2019). Overall, job interviews increase 44% in the treatment relative to control group, significant at the 10% level. The effects are concentrated among those who are occupationally narrow, where interviews double for those in the treatment relative to the control group. That means that the effects are concentrated among those who are arguably not looking broad enough on their own, likely due to their lack of knowledge of possible alternatives. Those that are initially occupationally broad without the intervention do not seem to be negatively affected by the advice.

It is particularly noteworthy that point estimates are the largest for the effects that can be directly traced on the job search platform. But searching on the job search

TABLE 3. Effect of intervention on job interviews.

	Number of interviews		
	(1) Lab	(2) Survey	(3) Total
Treatment	0.61 (0.79)	0.40* (0.27)	0.44* (0.28)
Treatment			
X occupationally broad	− 0.37 (0.43)	− 0.00 (0.28)	− 0.07 (0.24)
X occupationally narrow	1.13 (1.26)	0.86** (0.47)	1.03*** (0.55)
Model	Poisson	Poisson	Poisson
Observation weeks	1–10	1–10	1–10
N	540	466	464

Source: Belot, Kircher, and Muller (2019) Table 10. Numbers represent point estimates from a Poisson model, brackets provide standard errors, stars indicate 10% (*) and 5% (**) significance, and occupationally broad refers to the subject pool who had above average occupational breadth in the 3 weeks before the intervention while occupationally narrow refers to the complement.

platform does not generate the bulk of interviews, and those arise outside of the platform are recorded in the answers to a survey question. There appears to be positive spillovers, and the information learned on the platform is also used by job seekers in other search activities. Further slicing the data indicates that it is particularly the longer-term unemployed that benefit from the intervention. This might be because those with short unemployment duration have strong priors about suitable occupations, and are therefore not interested in alternatives. The paper offers a toy model to illustrate this point.

While the effects on the individuals are economically substantial, the paper lacks insights in two dimensions. First, the study was underpowered to discover effects on job finding. The confidence intervals for job finding include the large positive point estimates reported for interviews, but also very large negative estimates because precision is so low. This is a point of concern because job interviews in occupations that differ from the initial focus of the job seeker might not have the same chances of success. The study also does not investigate whether the benefits to those that are treated come at costs to others who are untreated.

5.2. A Medium-Scale Study of Occupational Advice to Long-Term Unemployed Job Seekers

The previous study used a low number of job seekers with low to medium unemployment duration and relied on self-reported data. Belot, Kircher, and Muller (2022a) implement automated advice for a middle-sized group of long-term unemployed job seekers in England in 2019–2020. This has several advantages. First,

as the discussion in Section 4 indicated and is explored in more depth in the next section, improvements for long-term job seekers are likely to have positive equilibrium effects even if they come at a cost to those with lower unemployment durations. So negative spillovers are much less likely to offset such effects. Moreover, effect sizes tend to increase by unemployment duration, as reported in the previous subsection and more generally documented in Card, Kluve, and Weber (2018). Therefore, the impact on job finding is potentially easier to detect in this sample. And finally, this group has received considerable policy attention, and low-cost improvements in their labor market prospects are of high policy relevance.

The analysis of the job finding prospects in this setting is aided by access to administrative rather than self-reported data. The results indicate large responses on the ability to attain stable jobs, of magnitudes similar to those for job interviews in the study in the previous subsection. The main difference is that this study now considers actual job findings, rather than just interviews. It also presents evidence that the advice causally changed the types of occupations where individuals obtain new jobs. Providing all details goes beyond the current paper, but is provided in Belot, Kircher, and Muller (2022a).

5.3. Relation to Studies That Provide Job Search Advice without Embedding It in an Online Job Search Platform

There exist other studies that have provided advice but do not integrate it into an online job search platform. A sequence of papers has considered in-person advice and positive effects on job finding have generally been reported. Yet such advice tends to cost in the order of several hundred to several thousand Euros and the costs often negate the benefits it offers (e.g., Crépon et al. 2013; Behaghel et al. 2014; Cottier et al. 2018). In contrast, online implementations such as those discussed in this paper are more narrowly focused (but potentially more replicable) and bear negligible marginal cost of adding additional job seekers. It might be surprising that online advice seems to have so large effect sizes that they can be compared to those of in-person advice, and one reason might be that it directly connects the advice to the vacancies in the market, which makes it very concrete and tangible. It is known from psychological research that examples have a large effect for keeping things in memory, and showing actual vacancies might achieve that.

Altmann et al. (2018) sent advice via regular mail to 40,000 job seekers in Germany. Provision of this type of advice is low-cost, but is not tailored to the individual job seekers. This study also did not focus on occupational advice. Rather, it provided general advice that might be useful to all job seekers, such as that proactive job search behavior pays off and that finding a job early might be useful to avoid negative duration dependence. They find small and statistically insignificant effects for the whole population, and significant positive effects for those at risk of long-term unemployment, though effect sizes are substantially lower than those reported in the studies above. This could be due to many features that differ between these studies, including the

nature and personalization of advice. But they also differ in how advice is linked to actual jobs in the market.

Dhia et al. (2022) analyze an intervention that send emails to 120,000 job seekers in France to promote a government website that offers a wide range of advice for job search, including advice akin to that in the preceding paragraph as well as some forms of occupational advice. The intervention induced 13% in the treatment group to sign up to the site, while sign-ups in the control group were negligible. They do not find any impact on employment, either in general or for subgroups. The authors speculate that this could be due to general equilibrium effects. But since treatment and control group interact in the same markets and the study only exposes around 8% of individuals in each market to advice, it is unlikely that equilibrium effects could undo individual effects of the size reported in the previous subsections.⁸ Several other explanations could account for the differences: Apart from the obvious point that the smaller studies could be an outlier, the type of advice and the complexity of advice is different in this study. In addition, the advice is not directly linked to vacancies and might therefore not be directly converted into actual changes in job search behavior by individuals.

Overall, it is clear that not every study that introduces low-cost advice achieves large employment gains. Open questions include: Which types of advice are most suitable to help individuals? How can this advice be represented without cognitively overloading the participants? Does it matter if the advice is directly linked to open vacancies? Can one structure the job search process on a website to make it easier to absorb the advice, to turn it into action, and for participants to learn from it so they can use the knowledge even during informal job search? And which types of advice have the lowest negative spill-overs on other individuals, that is, the lowest general equilibrium effects.

Answering some of these questions will require trial and error, and will likely explore the benefits of new methods such as machine learning. Researchers will also need to draw on the psychology of human learning, and they will need an understanding of information overload. But some of the answers might also be guided by economic theory. The following section aims to provide some initial insights. It focuses on the question where advice might be least costly in terms of general-equilibrium externalities. As these studies get rolled out on larger scales and evaluated by the research community, these considerations can provide some guidance where positive effects for the overall population might be likely to obtain. They do not provide an exhaustive list, and have to be taken with some caution as the assumptions that underly them about the transferability of skills are strong. But they formalize current folk intuition about the type of settings where larger roll-out seems promising.

8. The previous subsections treated a very small fraction of the population, and show large effects of treated versus untreated individuals. The study in Dhia et al. (2022) treats a larger share, but the equilibrium effects it induces impact both treated and non-treated individuals. If beneficial individual effects were strong but are counteracted by general equilibrium effects, then the general equilibrium effects would have to be large and they would have to negatively effect in particular those in the treatment group. If the equilibrium effects were uniformly distributed, the treated individuals would continue to do better than the non-treated individuals.

6. A Stylized Toy Model

Advice has an impact if it changes where job seekers search. In a rational model, if advice changes behavior, it helps the individual at least in expectation. But it might include externalities on others, and overall the participants in the market could be worse off.

These “general equilibrium” spillovers on others can be explored empirically in large-scale randomized trials. There are some established research designs that allow to uncover these (see most prominently Crépon et al. (2013) who explore this for in-person advice). A large player in the online job-search market could roll out advice at different intensities across regional markets to explore the general-equilibrium spillovers, and the low marginal cost of online advice should open this possibility in the near future. Yet, so far this has not been explored empirically in the academic research community.

In order to provide some guidance on the type of settings in which negative spillovers might be small, this section relies on insights from economics theory. It provides formal underpinnings to the discussion in Section 4, collecting in a stylized way the basic folk knowledge from labor macro models. While this does not make up for empirical work on this topic, it can shine the light on the types of environments that are most likely to bring beneficial results in the aggregate. The author believes that the following points are intuitive, and he is not aware of any succinct theoretical presentations of these basic insights, which is the aim of this section.

Tractability in this section is achieved by neglecting the exact specification of the priors and signals and updating of beliefs. Rather, it identifies changes in behavior that benefit individuals in a particular group (in which case arguably some signal structure can be found that induces some of these individuals to indeed change their behavior in this direction). It then asks whether a general change in behavior by members of this group improves welfare for the overall economy.

Consider an economy with a mass N of workers. Workers have different types N_t for $t \in \{1, 2, \dots, T\}$. Workers can search for jobs in different sectors $s \in \{1, \dots, S\}$ that form separated submarkets. Only unemployed workers search for jobs, and the search intensity (or effort) of workers of type t in sector s is denoted by α_{ts} . Per unit of search intensity, individuals differ in their search efficiency, which is denoted by a_{ts} . Unemployed workers get a flow utility b , which is the sum of their intrinsic value for leisure b_L and a government transfer b_G . Employed workers obtain a flow utility equal to their wage w_{ts} .

There are V_s firms in sector s . Firms have a linear production function, and a marginal worker of type t employed in a firm in sector s has constant marginal product $\tilde{y}_{ts}$. The fraction of this output that is given to the worker is determined by Nash bargaining, where firms have a bargaining weight γ . For notational convenience we denote $y_{ts} = \tilde{y}_{ts} - b$ the output of a worker net of the unemployment benefits that the worker would have received if unemployed. Firms constantly advertise that they are

hiring, which is modeled as an active measure of vacancies that equals the number of firms V_s at all times.⁹

The economy operates in continuous time. All agents discount the future at the rate r . Meetings in the market are generated according to a matching function in the Diamond–Mortensen–Pissarides tradition: The meeting rate in a submarket is given by the meeting function $m(u_s, V_s)$ that is strictly increasing in the efficiency-weighted number of unemployed u_s and the number of vacancies V_s . It is also assumed to be strictly concave, and features constant returns to scale. The meeting rate per efficiency unit of search effort for a worker in this submarket is given by $q_s := q(\theta_s) = q(u_s/V_s) = m(u_s, V_s)/u_s$. The meeting rate for a firm is given by $p(\theta_s) = p(u_s/V_s) = m(u_s, V_s)/V_s$. Denote the stock of unemployed of type t by U_t , and the total stock of unemployed is $U = \sum_t U_t$. Weighted by search efficiency and effort we have $u_{ts} = a_{ts}\alpha_{ts}U_t$, and $u_s = \sum_t u_{ts}$. Employed workers get separated from their firms at an exogenous rate δ and return to the pool of unemployed.

Value Functions. The flow value for an unemployed worker of type t is given in the usual continuous-time value function specification:

$$rW_t = b + \sum_s a_{ts}\alpha_{ts}q(\theta_s) \max\{R_{ts} - W_t, 0\}, \quad (1)$$

where the left hand side captures the flow value from unemployment, which equals on the right-hand side the flow utility from unemployment (b), plus that chance to meet a firm (at rate $a_{ts}\alpha_{ts}q(\theta_s)$) multiplied by the gain in value from employment ($R_{ts} - W_t$ if this is positive). This includes the value of being employed in sector s , which can analogously be specified as

$$rR_{ts} = w_{ts} + \delta[W_t - R_{ts}], \quad (2)$$

which captures on the right hand side the flow wage (w_{ts}) plus the rate of losing the job times the associated difference in value when being unemployed rather than employed.

From the firms' perspective, the value of employing a worker of type t in sector s is

$$rQ_{ts} = \tilde{y}_{ts} - w_{ts} - \delta Q_{ts}, \quad (3)$$

which captures the output minus the wage, plus the separation probability times the loss in future value. If a vacancy meets a worker but does not form a match, the continuation value is zero (as the firm posts new vacancies in the next instance independent of whether the current one is filled or not).

9. One could alternatively introduce a free entry condition, but since this work focuses mainly on the worker's side this is left for future research.

The surplus of a match is the sum of the value of the match between a firm and a worker minus the value from remaining unmatched:

$$S_{ts} = Q_{ts} + R_{ts} - W_t, \quad (4)$$

and matches form if and only if this is positive. Let I_{ts} be an indicator function that takes the value of 1 if $S_{ts} \geq 0$ and zero otherwise. The wage is then negotiated such that $\gamma S_{ts} = Q_{ts}$ and $(1 - \gamma)S_{ts} + W_t = R_{ts}$.

Flow Equations. In steady-state the inflow of workers of type t into jobs of type s has to equal the outflow, so that

$$a_{ts}\alpha_{ts}q(\theta_s)I_{ts}U_t = \delta E_{ts}, \quad (5)$$

where E_{ts} is the employment of workers of type t in firms in sector s and the right hand side captures the rate of separations of workers from this sector, while the left hand side captures the inflow of new employees in this sector as the number of unemployed times their rate of finding work in this sector. The sum of unemployed and employed of each type has to equal the total population of that type:

$$U_t + \sum_s E_{ts} = N_t. \quad (6)$$

Finally, the steady-state output of the economy is given by

$$\begin{aligned} O &= \sum_{ts} \tilde{y}_{ts} E_{ts} + b_L U, \\ &= \sum_{ts} y_{ts} E_{ts} - b_G U + Nb, \end{aligned} \quad (7)$$

which sums output across all employed agents as well as leisure utility for unemployed agents. Unemployment benefits that are transfers are dropped, because they have to be financed from the output that is produced. The second line is simply a different representation, and highlights that more unemployment is cost the government resources in terms of unemployment benefits. Clearly, the last term in the second line is an additive constant that can be dropped from the planner's objective without loss of insight. While the planner's objective is standard for models with transferable utility, it might be worth-while to reiterate that the planner's objective is only pareto-improving relative to other allocations if the planner can use transfers. In principle, some groups might be hurt as long as the sum in (7) improves, and for a pareto-improvement these would need to be compensated.

The following subsection provides examples of situations where advice is beneficial according to objective (7). Note that here we take the initial actions of the individuals (α_{ts}) and the allocation of vacancies (V_s) as given, and ask what would happen if we could alter them mechanically. We will make sure that the altered actions α'_{ts} are indeed an improvement for the individual, which would be a requirement in any model with actual Bayesian learning where advice can only help individuals (at least in expectation). The real question here is not how to alter behavior, but rather under which

conditions the improvements for the individual are not offset by negative impacts on others.

6.1. Equating Market Tightness

If markets differ by tightness, it is efficiency improving to move workers from markets with relatively few jobs to markets with relatively many jobs. To see this in a stylized setting, assume there are only two worker types ($T = 2$) and two submarkets ($S = 2$), and workers have equal productivity y and equal search efficiency $a = 1$ (i.e., $y_{ts} = y$ and $a_{ts} = 1$ for all t and s). Assume workers only search in submarkets that equal their type, so that $\alpha_{11} = 1$ and $\alpha_{22} = 1$ and $\alpha_{ts} = 0$ otherwise. This might simply reflect their previous affiliations. Assume that $N_1/V_1 > N_2/V_2$, so that there are more workers per vacancy in sector 1, and evidently it is harder to obtain jobs in sector 1 than in sector 2.

PROPOSITION 1. *If workers of type 1 change their search intensity to $\alpha'_{11} = \alpha_{11} - \varepsilon$ and $\alpha'_{12} = \varepsilon$, this strictly improves the planner's objective for ε small enough. The planner's objective is maximized when $N_1\alpha'_{11}/V_1 = (N_1\alpha'_{12} + N_2\alpha_{22})/V_2$.*

Proof. See Appendix A. □

From a macro-labor perspective, this insight is obvious and follows from the concavity of the matching function $m(\cdot, \cdot)$. All else equal, this concavity implies that more matches are created the more equal the market tightness θ across sectors. This is intuitive: If nearly all workers are searching in one sector and nearly all vacancies are in the other sector, not many matches are created. By shifting workers to the sector with more vacancies, the additional congestion in the sector with many vacancies is minimal while the relief in the sector with many workers is substantial. Matches are maximized when market tightness θ is equalized across sectors, which is the condition in the proposition. A small ε implies more equal market tightness.

Clearly, redirecting individuals of type 1 to sector 2 is also in their own interest, as it trivially increases their job-finding rates by searching more in the “easy” sector. The question that this section addresses is that the negative consequences on workers of type 2 are lower than the benefits for workers of type 1, and so the individual benefit actually translates into an aggregate benefit. One can also easily show that in this setting the following holds: The larger the individual benefit, the larger the aggregate benefit. Individual benefit arises only if market tightness is rather unequal, and this is exactly the condition for the aggregate benefit to be large.

Some caveats are in order, both here and for all subsequent illustrations: This example was built around the idea that workers of the first type are equally able to do the work in the second sector. By continuity, the improvement from assigning more type-1 workers to sector 2 also holds when y_{12} is slightly smaller than y , so if they are slightly less productive in the second sector. But if their productivity is much lower in the second sector, evidently this insight is no longer true. Take the

extreme case where $y_{12} = 0$. Inducing job search effort of type-1 workers in sector 2 negatively affects type 2 workers, but also negatively affects sector 2 firms who now run into unproductive workers. It also hurts the type-1 workers themselves, as they spend energy on unproductive jobs that they would not obtain, as well as type-1 firms that now meet suitable workers less often.

6.2. Inducing More Search

Take the same setup as the one introduced at the beginning of the previous subsection, with the only difference that now $N_1/V_1 = N_2/V_2$. This means that market tightnesses are equalized across sectors. But instead of considering only their own sector, assume now that information induces workers to search in both sectors, so it changes their search to $\alpha'_{ts} = 1$ for all t and s .

PROPOSITION 2. *Search strategy α'_{ts} strictly improves the planner's objective relative to search strategy α_{ts} .*

Proof. See Appendix A. □

The reason behind this result is again simple: As workers realize that they can also search and be productive in the other sector, the economy effectively converts from a representative-type economy with a search effort of 1 into a representative-type economy with a search effort of 2. Higher aggregate search effort improves the number of overall matches.

It might be noteworthy to stress that general equilibrium effects are still present: if everyone else searches with search intensity α_{ts} that focuses only on one sector, and only a finite number of individuals are induced to adopt the search strategy α'_{ts} that searches across both sectors, the effect on the finite sample will be larger than the effect if everyone adopts strategy α'_{ts} . The reason is obvious: When everyone else adopts strategy α_{ts} , then the tightness θ_s in the markets is rather low. This means that the chances to find jobs in any market is high, and the additional benefit from searching in the second market is high. Obviously, when everyone adopts the strategy α'_{ts} , then $\theta'_s > \theta_s$ and it is harder to find a job in any given market.

Also, the caveat of equal productivity still applies, and it becomes obvious that the setup abstracted from the costs of more search effort. Again by continuity, the argument continues to hold even if y_{12} and y_{21} are slightly below y , and if the search effort has some slight costs. But it can reverse if search effort is extremely costly or productivity in the other sector is particularly low.

6.3. Inducing Individually Better Search

Section 6.1 was intended to represent a setting where there are entire submarkets with much more job opportunities than other submarkets and some individuals are not exploring this. Alternatively, one can think of settings in which some individuals

currently search in markets where their own productivity or ability to find jobs is low, even if the market overall has relatively many vacancies. In reality, this might happen in particular to unemployed job seekers who were previously employed for a long time in one occupation, and who think that they are still able to find a job in this occupation, even though the requirements of or the strategies to obtain such jobs have moved on. But it could also happen to unemployed workers without sufficient experience to calibrate their job search strategy to the requirements of the jobs.

For exposition, continue with the setting with $T = 2$ types of workers and $S = 2$ sectors. Productivity is constant at y (i.e., $y_{ts} = y$ for all t, s), efficiency-weighted search effort is constant ($a_{ts}\alpha_{ts} = \lambda$ for all t, s) and vacancies and tightness are equalized ($V_1 = V_2$). But worker type 1 has lower search efficiency in sector 1 (i.e., $a_{11} < a_{12}$) and nevertheless maintains higher search effort in this sector (i.e., $\alpha_{11} > \alpha_{12}$), possibly reflecting past affiliations. The next proposition highlights that the planner would rather shift this type's search effort more toward sector 2, where his search efficiency is higher, despite potential negative consequences on the other worker type.

PROPOSITION 3. *All else equal, the planner's objective strictly increases under alternative strategy $\alpha'_{11} = \alpha_{11} - \varepsilon$ and $\alpha'_{12} = \alpha_{12} + \varepsilon$ for ε sufficiently small.*

Proof. See Appendix A. □

The reason for the improvement is obvious: By shifting search effort from sector 1 to sector 2, the individuals switch their focus to the sector where they have a structural advantage in getting a job. This helps these workers.

Note that the shift increases efficiency-weighted search effort overall. Therefore, it is relevant that the productivity of type-1 workers is not much lower than the productivity of other workers that search in these sectors, since some matches might be relocated to type-1 workers.

6.4. Helping those at Risk of Long-Term Unemployment

Another intervention that is intuitively helpful is to provide jobs to those at the risk of long-term unemployment, even if this means taking jobs away from others and does not raise the overall employment rate. The reason is that others have a much quicker time to receive alternative offers, so overall there are savings in government subsidies.

To see this, consider now a market where some individuals, say those of type 1, have lower matching efficiency than others, say those of type 2 (i.e., $a_{1j} = a_1 < a_{2j} = a_2$). We assume that the differences in matching efficiency are not too large,¹⁰ and there are three markets: Group 1 searches only in markets 1 and 3, and group 2 searches only in markets 2 and 3, with search intensities $\alpha_{11} = \alpha_{12} = \alpha$ and $\alpha_{22} = \alpha_{23} = \alpha$.

10. A sufficient condition that makes this precise is provided in Assumption A.1 in the Appendix A.

Productivity in all markets is y , and both groups have equal sizes. Assume the number of vacancies across the three markets are such that market tightnesses in steady state are initially equalized. Faced with equal market tightness everywhere, clearly, type-1 workers have higher unemployment and any unemployed member in that group has higher chances to be long-term unemployed.

Now consider a website that redirects traffic such that a small number of vacancies ε are shown more to group 1, modeled as a shift from market 3 to market 1. That is, assume ε such that $V'_3 = V_3 - \varepsilon$, $V'_1 = V_1 + \varepsilon$, and $V'_2 = V_2$. All else is unchanged.

PROPOSITION 4. *For sufficiently small ε , the planner's objective improves, even though the number of matched agents in equilibrium falls.*

Proof. See Appendix A. □

The reason that the number of matches falls is that market tightness is no longer equalized across sectors. This is made up for by shifting job findings to the low-finding agents, while the high-finding agents can secure more easily additional alternatives. The effect on the number of matches is second-order because of equalized market tightness, while the change in unemployment benefits has first-order effects.

Note again that the choice to model differences through search efficiency rather than through productivity is not innocuous: It is important that such interventions are targeted to sectors where productivity of long-term job seekers does not significantly lag behind that of other job seekers. If productivity differentials are large, shifting vacancies towards low-productivity employees does provide aggregate benefits.

6.5. Model Discussion

The model is kept simple to provide some basic insights. The conditions underlying the propositions are strong, including those on the base allocation of search effort and vacancies upon which the planner wants to improve. Deeper theoretical work might be needed to explore the exact boundaries of these insights. The subsections before already highlighted some especially relevant caveats. In particular, moving the search activities of individuals to other occupations/sectors is shown to be beneficial if the productivity there is not substantially lower than in their original sector or than that of those who usually seek work in this sector. Evidently, advice that induces individuals to search where their own productivity is low is unlikely to generate the same positive impact on the planner's objective.

One other shortcoming of the model above is that the change in the actions is imposed mechanically. More appropriate would be a learning model where individuals take particular actions based on their prior beliefs about the benefits, but additional signals move these beliefs and, therefore, the agents' optimal actions. Standard models of this kind have the feature that additional information in expectation helps the receiver (as she could always choose to disregard the information). To that end, our analysis in the previous subsections choose settings in which the change in actions indeed benefit

the affected individuals privately. This makes it conceivable that the initial actions and the changes could be micro-founded in a learning model.

Belot, Kircher, and Muller (2019) lay out one such learning model and investigate its implications. They in particular point out why information provision might not affect job seekers early on during unemployment, but only after later some months in unemployment. This happens in their illustration if job seekers initially have a good prior where to search, and so focus their search on those target occupations. They disregard the information on other occupations and learn mainly in their target occupations. Job seekers that are unsuccessful for several months in their target occupations downgrade their belief about the chances in that occupation and consider alternatives. At that point, they rely on advice on which alternatives to pursue. Belot, Kircher, and Muller (2019) outlined this but do not investigate general equilibrium spillovers.

A particularly interesting feature about modeling beliefs more carefully is that useful advice on a platform can alter the behavior of an individual also on other platforms or in informal job search through friends and family. Take the example of occupational advice. If an individual obtains such advice on one platform, she might also use it when she searches elsewhere or when she contacts her social network for referrals. The empirical results in Belot, Kircher, and Muller (2019) suggest that this is indeed a consequence of occupational advice in their setting.

Overall, this section presents in a stylized setting where interventions might be fruitful in the sense that their individual benefits outweigh the negative externalities on others. This arises when it induces individuals to leave occupations with particularly few vacancies per job seeker, to search more, to search in sectors where their ability to find jobs is higher, or if it helps long-term unemployed who otherwise face protracted unemployment. Studies that explore this empirically in the context of online advice are still missing, as are theories that derive these insights more generally and in a full-fledged learning framework. This theoretical framework provides hints where future empirical studies might want to focus to find positive aggregate effects.

7. Conclusion

Formal job search has moved from offline to online during the last 20 years, with platforms now acting as large intermediators in this market. This provides a setup that generates exciting new data about the process by which workers and firms find each other. Currently, this allows researchers to generate new insights into both the preferences of firms as stated in their vacancy announcements, as well as the search behavior of workers exhibited in their queries, views, and applications. As these platforms evolve, they aim to make the entire process of match formation easier. It is likely that this will generate more data on how firms react to the job applications, and who the final hire is. While such data is currently still rare, it will provide a nearly holistic picture of the interactions that take place in the match formation.

The digital nature of the intermediating job search platforms makes it also relatively low-cost to change them and to assess the consequences in a randomized controlled trial. An obvious domain where efforts are made both on the academic and commercial front is to provide recommendations to job seekers (and also to recruiters on the other side of the market). Likely the effects of advice lie in the details: which kind of advice, how it is presented, and how it is interfaced with the rest of the job search process. There are early studies that show promise, but more research is needed.

Still, studies that assess the general equilibrium effects of rolling out such advice at large scale are lacking. Such studies could soon be feasible, though, since some private or public players in the online market cover a non-trivial part of it in some countries. This paper presented in an illustrative model some scenarios of occupational advice in which the direct effects of the advice are likely to dominate the negative general equilibrium spillovers on others. Large-scale rollouts might focus on such scenarios to have larger chances to detect positive overall effects.

Appendix A: Proofs of Propositions

A.1. Proposition 1: Equating Market Tightness

Suppose $N_1/V_1 > N_2/V_2$; $y_{t,s} = y$ for $t, s \in \{1, 2\}$; $a_{t,s} = 1$ for $t, s \in \{1, 2\}$; and $\alpha_{t,s} = 1$ if $t = s$ and 0 otherwise.

PROPOSITION 1. *If workers of type 1 change their search intensity to $\alpha'_{11} = \alpha_{11} - \varepsilon$ and $\alpha'_{12} = \varepsilon$, this strictly improves the planner's objective for ε small enough. It is maximized when $N_1\alpha'_{11}/V_1 = (N_1\alpha'_{12} + N_2\alpha_{22})/V_2$.*

Proof. Let us prove separately the two parts of the proposition. Steps of the proof for part 1 only:

1. rewrite the planner's objective in ε ;
2. define unemployment and tightness in ε ;
3. compute the unemployment and tightness derivative in ε ;
4. evaluate the derivatives at $\varepsilon = 0$;
5. show the planner's objective strictly improves for $\varepsilon > 0$ sufficiently small.

1. *Rewrite the planner's objective in ε .* The planners objective is

$$O = y(E_{11} + E_{12} + E_{22}) - b_G(U_1 + U_2).$$

Use the flow equation $E_{ts} = \alpha_{ts}q(\theta_s)U_s/\delta$ and the market clearing condition $N_t = E_{t,1} + E_{t,2} + U_s$ to obtain

$$O(\varepsilon) = y(N_1 + N_2) - (y + b_G)U(\varepsilon),$$

$$U(\varepsilon) = U_1(\varepsilon) + U_2(\varepsilon),$$

so that the planner's objective is to reduce aggregate unemployment $U(\varepsilon)$ across the worker types.

2. *Define unemployment and tightness in ε .* Use $\alpha_{1,1} = (1 - \varepsilon)$, $\alpha_{1,2} = \varepsilon$ and $\alpha_{2,2} = 1$ to write

$$U_1(\varepsilon) = \frac{\delta N_1}{\delta + (1 - \varepsilon)q(\theta_1(\varepsilon)) + \varepsilon q(\theta_2(\varepsilon))},$$

$$U_2(\varepsilon) = \frac{\delta N_2}{\delta + q(\theta_2(\varepsilon))},$$

$$\theta_1(\varepsilon) = \frac{(1 - \varepsilon)U_1(\varepsilon)}{V_1},$$

$$\theta_2(\varepsilon) = \frac{\varepsilon U_1(\varepsilon) + U_2(\varepsilon)}{V_2}.$$

3. *Compute the unemployment and tightness derivative in ε .*

$$U'_1(\varepsilon) = - \frac{\delta N_1[-q(\theta_1(\varepsilon)) + (1 - \varepsilon)q'(\theta_1(\varepsilon))\theta'_1(\varepsilon) + q(\theta_2(\varepsilon)) + \varepsilon q'(\theta_2(\varepsilon))\theta'_2(\varepsilon)]}{[\delta + (1 - \varepsilon)q(\theta_1(\varepsilon)) + \varepsilon q(\theta_2(\varepsilon))]^2},$$

$$U'_2(\varepsilon) = - \frac{\delta N_2 q'(\theta_2(\varepsilon))\theta'_2(\varepsilon)}{[\delta + q(\theta_2(\varepsilon))]^2},$$

$$\theta'_1(\varepsilon) = \frac{(1 - \varepsilon)U'_1(\varepsilon) - U_1(\varepsilon)}{V_1},$$

$$\theta'_2(\varepsilon) = \frac{U_1(\varepsilon) + \varepsilon U'_1(\varepsilon) + U'_2(\varepsilon)}{V_2}.$$

4. *Evaluate the derivatives at $\varepsilon = 0$.* Simplify notation as follows:

$$U_1 = U_1(\varepsilon), \quad U_2 = U_2(\varepsilon), \quad \theta_1 = \theta_1(\varepsilon), \quad \theta_2 = \theta_2(\varepsilon),$$

$$U'_1 = U'_1(\varepsilon), \quad U'_2 = U'_2(\varepsilon), \quad \theta'_1 = \theta'_1(\varepsilon), \quad \theta'_2 = \theta'_2(\varepsilon).$$

Evaluate the derivatives from point 3. at $\varepsilon = 0$:

$$U'_1 = - \frac{q(\theta_2) - q(\theta_1) + q'(\theta_1)\theta'_1}{\delta + q(\theta_1)} U_1,$$

$$U'_2 = - \frac{q'(\theta_2)\theta'_2}{\delta + q(\theta_2)} U_2,$$

$$\theta'_1 = \frac{U'_1}{V_1} - \theta_1,$$

$$\theta'_2 = \frac{U_1 + U'_2}{V_2}.$$

Substitute the values of θ'_1 and θ'_2 into U'_1 and U'_2 and rearrange to obtain

$$U'_1 = \frac{-q(\theta_2) + q(\theta_1) + q'(\theta_1)\theta_1}{\delta + q(\theta_1) + q'(\theta_1)\theta_1} U_1,$$

$$U'_2 = \frac{-q'(\theta_2)\theta_2}{\delta + q(\theta_2) + q'(\theta_2)\theta_2} U_1.$$

5. Show the planner's objective strictly improves for $\varepsilon > 0$ sufficiently small. We need to show that $U'_1 + U'_2 < 0$. Notice

$$U'_2 < \frac{-q'(\theta_2)\theta_2}{\delta + q(\theta_1) + q'(\theta_1)\theta_1} U_1,$$

since $q(\theta) + q'(\theta)\theta$ is decreasing in θ and $\theta_2 < \theta_1$. To see $q(\theta) + q'(\theta)\theta$ is decreasing observe that it holds only iff

$$\theta q''(\theta) + 2q'(\theta) < 0.$$

Note that $q(\theta) = m(1, 1/\theta)$, so that $q'(\theta) = -m_V/\theta^2$ and $q''(\theta) = m_{VV}/\theta^4 + 2m_V/\theta^3$. Back in the previous condition this yields

$$\begin{aligned} m_{VV}/\theta^3 + 2m_V/\theta^2 - 2m_V/\theta^2 &< 0 \\ \Leftrightarrow m_{VV}/\theta^3 &< 0, \end{aligned}$$

which holds by concavity of m . Therefore, we can write

$$U'_1 + U'_2 < \frac{-[q(\theta_2) + q'(\theta_2)\theta_2] + [q(\theta_1) + q'(\theta_1)\theta_1]}{\delta + q(\theta_1) + q'(\theta_1)\theta_1} U_1 < 0,$$

where the last inequality holds because both the numerator and denominator are positive (the former due to $q(\theta) + q'(\theta)\theta$ being decreasing in θ and the latter since $q(\theta) + q'(\theta)\theta > 0$). This concludes the proof of the first part of the proposition.

The steps to prove the second part of the proposition are:

1. compute the change in effort for type 1 worker;
2. show the new effort composition minimizes unemployment.

1. *Compute the change in effort for type 1 worker.* To simplify notation, fix $\alpha_{22} = \alpha_{11} = \alpha'_{11} + \alpha'_{12} = 1$ and define $\hat{\varepsilon}$ so that $\alpha'_{11} = (1 - \hat{\varepsilon})\alpha_{11}$. We want to show that $\hat{\varepsilon}$ fulfilling

$$\frac{(1 - \hat{\varepsilon}N_1)}{V_1} = \frac{\hat{\varepsilon}N_1 + N_2}{V_2},$$

$$\hat{\varepsilon} = \frac{N_1V_2 - N_2V_1}{N_1(V_1 + V_2)}.$$

maximizes the planner's objective under the constrain $\alpha'_{11} + \alpha'_{12} = 1$. To show this, we are going to show that $\hat{\varepsilon}$ minimizes unemployment, i.e. the planner's equivalent objective as from above. Notice $\varepsilon \in (0, 1)$ as $N_1/V_1 > N_2/V_2$.

2. *Show the new effort composition minimizes unemployment.* Following the notation from the first part of the proof, we want to show that $U'(\hat{\varepsilon}) = U'_1(\hat{\varepsilon}) + U'_2(\hat{\varepsilon}) = 0$.¹¹

Use the expression derived in part 1 to write

$$\begin{aligned} U'(\hat{\varepsilon}) &= \frac{-\delta N_1[(1 - \hat{\varepsilon})q'(\theta_1(\hat{\varepsilon}))\theta'_1(\hat{\varepsilon}) + \hat{\varepsilon}q'(\theta_2(\hat{\varepsilon}))\theta'_2(\hat{\varepsilon})]}{\delta + q(\theta_1(\hat{\varepsilon}))} + \frac{-\delta N_2q'(\theta'_2(\hat{\varepsilon}))\theta'_2(\hat{\varepsilon})}{\delta + q(\theta_2(\hat{\varepsilon}))} \\ &= \frac{-\delta N_1[(1 - \hat{\varepsilon})q'(\theta_1(\hat{\varepsilon}))\theta'_1(\hat{\varepsilon}) + \hat{\varepsilon}q'(\theta_2(\hat{\varepsilon}))\theta'_2(\hat{\varepsilon})] - \delta N_2q'(\theta'_2(\hat{\varepsilon}))\theta'_2(\hat{\varepsilon})}{\delta + q(\theta_1(\hat{\varepsilon}))} \\ &= \frac{\delta q'(\theta_1(\varepsilon))}{\delta + q(\theta_1(\hat{\varepsilon}))} [(1 - \hat{\varepsilon})N_1\theta'_1(\hat{\varepsilon}) + (N_1\hat{\varepsilon} + N_2)\theta'_2(\hat{\varepsilon})], \end{aligned}$$

since market tightness at $\varepsilon = \hat{\varepsilon}$ is equal in sector 1 and 2 that implies $q(\theta_1(\hat{\varepsilon})) = q(\theta_2(\hat{\varepsilon}))$, and $q'(\theta_1(\hat{\varepsilon})) = q'(\theta_2(\hat{\varepsilon}))$. Evaluate $\theta'_1(\varepsilon)$ and $\theta'_2(\varepsilon)$ at $\varepsilon = \hat{\varepsilon}$, and use the definition of $\hat{\varepsilon}$ from point 1 above to obtain

$$\begin{aligned} U'(\hat{\varepsilon}) &= \frac{\delta q'(\theta_1(\varepsilon))}{\delta + q(\theta_1(\hat{\varepsilon}))} \left[\frac{(1 - \hat{\varepsilon})N_1}{V_1} [(1 - \hat{\varepsilon})U'_1(\hat{\varepsilon}) - U_1(\hat{\varepsilon})] \right. \\ &\quad \left. + \frac{N_1\hat{\varepsilon} + N_2}{V_2} [U_1(\hat{\varepsilon}) + \hat{\varepsilon}U'_1 + U'_2(\hat{\varepsilon})] \right] \\ &= \frac{\delta q'(\theta_1(\varepsilon))}{\delta + q(\theta_1(\hat{\varepsilon}))} \frac{(1 - \hat{\varepsilon})N_1}{V_1} [(1 - \hat{\varepsilon})U'_1(\hat{\varepsilon}) - U_1(\hat{\varepsilon}) + U_1(\hat{\varepsilon}) + \hat{\varepsilon}U'_1 + U'_2(\hat{\varepsilon})] \\ &= \frac{\delta q'(\theta_1(\varepsilon))}{\delta + q(\theta_1(\hat{\varepsilon}))} \frac{(1 - \hat{\varepsilon})N_1}{V_1} U'(\hat{\varepsilon}), \end{aligned}$$

□

which to hold true implies $U'(\hat{\varepsilon}) = 0$.

11. Additionally, one would want to show that $\hat{\varepsilon}$ actually identifies a minimum by fulfilling the second order condition, that is $U''(\hat{\varepsilon}) > 0$.

A.2. Proposition 2: Inducing more Search

Suppose $N_1/V_1 = N_2/V_2$; $y_{t,s} = y$ for $t, s \in \{1, 2\}$; $a_{t,s} = 1$ for $t, s \in \{1, 2\}$; and $\alpha_{t,s} = 1$ if $t = s$ and 0 otherwise.

PROPOSITION 2. *Search strategy $\alpha'_{ts} = 1 \forall s, t$ strictly improves the planner's objective relative to search strategy α_{ts} .*

Proof. Assume $N_2 = \gamma N_1$ with $\gamma > 0$. Since $N_1/V_1 = N_2/V_2$ we have $V_2 = \gamma V_1$. As productivity y is identical across worker types and sectors, the planner's problem improves when overall unemployment is lowered (see discussion in proof of proposition 1). In formulas, we want to show $U_1(\alpha) + U_2(\alpha) > U_1(\alpha') + U_2(\alpha')$.

Since $N_1/V_1 = N_2/V_2$, we have $\theta_1(\alpha) = \theta_2(\alpha) = \theta(\alpha)$. Using the flow equations and market clearing conditions, we can write

$$U_1(\alpha) + U_2(\alpha) = \frac{(1 + \gamma)\delta N_1}{\delta + q(\theta(\alpha))},$$

$$U_1(\alpha') + U_2(\alpha') = \frac{(1 + \gamma)\delta N_1}{\delta + q(\theta_1(\alpha')) + q(\theta_2(\alpha'))}.$$

It follows that $U_1(\alpha) + U_2(\alpha) > U_1(\alpha') + U_2(\alpha')$ iff

$$q(\theta_1(\alpha')) + q(\theta_2(\alpha')) > q(\theta(\alpha)),$$

where market tightness values are

$$\theta(\alpha) = \frac{\delta}{\delta + q(\theta(\alpha))} \frac{N_1}{V_1},$$

$$\theta_1(\alpha') = (1 + \gamma) \frac{\delta}{\delta + q(\theta_1(\alpha')) + q(\theta_2(\alpha'))} \frac{N_1}{V_1},$$

$$\theta_2(\alpha') = \frac{1 + \gamma}{\gamma} \frac{\delta}{\delta + q(\theta_1(\alpha')) + q(\theta_2(\alpha'))} \frac{N_1}{V_1}.$$

Assume, by way of contradiction, that $q(\theta_1(\alpha')) + q(\theta_2(\alpha')) \leq q(\theta(\alpha))$. It follows that

$$\theta_1(\alpha') < (1 + \gamma)\theta(\alpha),$$

$$\theta_2(\alpha') < (1 + \gamma)/\gamma \theta(\alpha).$$

Since $q'(\theta) < 0$, we can write

$$\begin{aligned} q(\theta_1(\alpha')) + q(\theta_2(\alpha')) &> q((1+\gamma)\theta(\alpha)) + q((1+\gamma)/\gamma \theta(\alpha)) \\ &= m\left(1, \frac{1}{(1+\gamma)\theta(\alpha)}\right) + m\left(1, \frac{\gamma}{(1+\gamma)\theta(\alpha)}\right) \\ &> m\left(1, \frac{1}{\theta(\alpha)}\right) \\ &= q(\theta(\alpha)), \end{aligned}$$

where the last inequality follows from the concavity of $m(u, v)$ in v . The last expression is a contradiction, so it must be $q(\theta_1(\alpha')) + q(\theta_2(\alpha')) > q(\theta(\alpha))$, i.e. $U_1(\alpha) + U_2(\alpha) > U_1(\alpha') + U_2(\alpha')$.

A.3. Proposition 3: Inducing Individually Better Search

Assume that worker type 1 has equal productivity in sectors 1 and 2 (i.e., $y_{11} = y_{12} = y$) but has lower search efficiency in sector 1 (i.e., $a_{12} > a_{11}$). But, because of previous affiliations, he spends too much time searching in sector 1 (i.e., $a_{11}\alpha_{11} = a_{12}\alpha_{12} = \lambda > 0$). Assume sectors 1 and 2 have equal number of firms (i.e., $V_1 = V_2 > 0$) and assume all type 2 workers spend equal efficiency units of search across these sectors, and they generate the same output as type-1 workers (i.e., $a_{21}\alpha_{21} = a_{22}\alpha_{22} = \lambda$ and $y_{21} = y_{22} = y$).

PROPOSITION 3. *All else equal the planner's objective strictly increases under alternative strategy $\alpha'_{11} = \alpha_{11} - \varepsilon$ and $\alpha'_{12} = \alpha_{12} + \varepsilon$ for ε sufficiently small.*

Proof. The planner's problem as a function of ε is

$$\begin{aligned} O(\varepsilon) &= y \sum_{ts} E_{ts}(\varepsilon) - b_G \sum_t U_t(\varepsilon) \\ &= y(N_1 + N_2) - [y + b_G](U_1(\varepsilon) + U_2(\varepsilon)) \end{aligned}$$

Below, we are going to show $U' = U'_1(0) + U'_2 < 0$.

Equilibrium unemployment and market tightness as a function of ε are

$$\begin{aligned} U_1(\varepsilon) &= \frac{\delta N_1}{\delta + \lambda[q(\theta_1(\varepsilon)) + q(\theta_2(\varepsilon))] + \varepsilon[-q(\theta_1(\varepsilon))a_{11} + q(\theta_2(\varepsilon))a_{12}]}, \\ U_2(\varepsilon) &= \frac{\delta N_2}{\delta + \lambda[q(\theta_1(\varepsilon)) + q(\theta_2(\varepsilon))]} \\ \theta_1 &= \frac{\lambda[U_1(\varepsilon) + U_2(\varepsilon)] - \varepsilon a_{11} U_1(\varepsilon)}{V_1} \\ \theta_2 &= \frac{\lambda[U_1(\varepsilon) + U_2(\varepsilon)] + \varepsilon a_{12} U_1(\varepsilon)}{V_1} \end{aligned}$$

Notice that market tightness is equal across sectors for $\varepsilon = 0$; that is $\theta := \theta_1(0) = \theta_2(0)$. Take the derivative in ε for each of the expressions above, and evaluate at $\varepsilon = 0$, to obtain (omit the derivative argument)

$$\begin{aligned} U'_1 &= \frac{-U_1}{\delta + 2\lambda q(\theta)} [\lambda q'(\theta)(\theta'_1 + \theta'_2) + q(\theta)(a_{12} - a_{11})], \\ U'_2 &= \frac{-U_2}{\delta + 2\lambda q(\theta)} [\lambda q'(\theta)(\theta'_1 + \theta'_2)], \\ \theta'_1 &= \frac{\lambda(U'_1 + U'_2)}{V_1} - \frac{a_{11}U_1}{V_1}, \\ \theta'_2 &= \frac{\lambda(U'_1 + U'_2)}{V_1} + \frac{a_{12}U_1}{V_1}. \end{aligned}$$

Plug market tightness derivatives into unemployment derivatives and rearrange

$$\begin{aligned} U'_1 &= \frac{-U_1}{\delta + 2\lambda q(\theta)} \left[q(\theta)(a_{12} - a_{11}) + 2\lambda q'(\theta) \frac{\lambda(U'_1 + U'_2)}{V_1} \right], \\ U'_2 &= \frac{-U_2}{\delta + 2\lambda q(\theta)} \left[\frac{\lambda U_1}{V_1} q'(\theta)(a_{12} - a_{11}) + 2\lambda q'(\theta) \frac{\lambda(U'_1 + U'_2)}{V_1} \right]. \end{aligned}$$

Add up unemployment derivatives, and rearrange, so to obtain

$$(U'_1 + U'_2) \left[1 + \frac{2\lambda q'(\theta)\theta}{\delta + 2\lambda q(\theta)} \right] = \frac{-U_1(a_{12} - a_{11})}{\delta + 2\lambda q(\theta)} \left[q(\theta) + q'(\theta) \frac{\lambda U_2}{V_1} \right],$$

which implies $(U'_1 + U'_2) < 0$. To get to this result, notice the following two facts. First, the left hand side is positive as $q(\theta) + q'(\theta)\theta > 0$ implies

$$\frac{2\lambda q'(\theta)\theta}{\delta + 2\lambda q(\theta)} \in (-1, 0).$$

Second, the right hand side is negative since $a_{12} - a_{11} > 0$ and $q(\theta) + q'(\theta) \frac{\lambda U_2}{V_1} > q(\theta) + q'(\theta)\theta > 0$.

We showed $(U'_1 + U'_2) < 0$, which implies the planner's objective increases for $\varepsilon > 0$ and small enough.

A.4. Proposition 4: Helping Those at Risk of Long-Term Unemployment

Assume $t \in \{1, 2\}$ and $s \in \{1, 2, 3\}$; productivity is $\forall s, t : y_{t,s} = y$; search efficiency is $\forall s : a_{1,s} = a_1 < a_2 = a_{2,s}$; and search effort is always α except $\alpha_{13} = \alpha_{21} = 0$. Vacancies V_s are such that market tightness $\forall s : \theta_s = \theta$. Suppose there is a policy intervention affecting vacancies $V'_1 = V_1 + \varepsilon$ and $V'_3 = V_3 - \varepsilon$, while the rest is unchanged.

ASSUMPTION A.1. Define $A := q(\theta) + q'(\theta)\theta$, then $a_1 < a_2 < a_1(1 + A) + 1$.¹²

PROPOSITION 4. *Under Assumption A.1, for sufficiently small ε the planners objective improves, even though the number of matched agents in equilibrium falls.*

Proof. Define $O(\varepsilon)$ and $E(\varepsilon)$ and $U(\varepsilon)$, respectively, as the planner's objective, employment and unemployment functions with ε as argument. We want to show $O'(0) > 0$.

We can use the market clearing condition $N_1 + N_2 = E(\varepsilon) + U(\varepsilon)$ to rewrite the planner's objective:

$$\begin{aligned} O(\varepsilon) &= yE(\varepsilon) - b_G U(\varepsilon) \\ &= [y + b_G]E(\varepsilon) - b_G[N_1 + N_2], \end{aligned}$$

so that $O'(\varepsilon) > 0$ iff $E'(\varepsilon)$.

For every submarket, write the market tightness θ_s and efficiency-weighted mass of unemployed u_s as a function of ε . We can write the following expression for employment:

$$E(\varepsilon) = \sum_s q(\theta_s(\varepsilon))u_s(\varepsilon),$$

$$\text{where } \theta_s(\varepsilon) = \frac{u_s(\varepsilon)}{V_s(\varepsilon)} \quad \text{and} \quad u_s(\varepsilon) = \sum_t \alpha_{ts} a_t U_t(\varepsilon).$$

Below, we are going to find two expressions for $E'(0)$ that we will later compare:

1. Differentiate $E(\varepsilon)$ and evaluate at $\varepsilon = 0$ to obtain (omit the argument when it is $\varepsilon = 0$):

$$E' = \sum_s q'(\theta)\theta'_s u_s + q(\theta)u'_s,$$

where expressions θ'_s and u'_s take values as follows:

$$\begin{aligned} \theta'_1 &= \frac{u'_1 - \theta}{V_1}, & \theta'_2 &= \frac{u'_2}{V_2}, & \theta'_3 &= \frac{u'_3 + \theta}{V_3}, \\ u'_1 &= \alpha a_1 U'_1, & u'_2 &= \alpha a_2 U'_2, & u'_3 &= \alpha(a_1 U'_1 + a_2 U'_2). \end{aligned}$$

Substitute terms and simplify to obtain

$$E' = 2\alpha[q'(\theta)\theta + q(\theta)](a_1 U'_1 + a_2 U'_2).$$

2. Differentiate $E(\varepsilon) = \sum_t N_t - U(\varepsilon)$ and evaluate at $\varepsilon = 0$ to obtain

$$E' = -U' = -(U'_1 + U'_2).$$

12. Recall $A > 0$ since the matching function had decreasing returns to scale.

Compare the two expressions for E' and solve for U'_1

$$U'_1 = -\frac{2\alpha[q'(\theta)\theta + q(\theta)]a_2 + 1}{2\alpha[q'(\theta)\theta + q(\theta)]a_1 + 1}U'_2 = -\frac{2\alpha Aa_2 + 1}{2\alpha Aa_1 + 1}U'_2.$$

Notice U'_1 and U'_2 have opposite signs as the fraction term is positive (recall that $A > 0$). Also $|U'_1| > |U'_2|$ as $a_2 > a_1$. Recall from the expression of E' in 1, that the sign of the employment derivative is the same of $a_1U'_1 + a_2U'_2$. Use the relation between unemployment derivatives to obtain the following

$$a_1U'_1 + a_2U'_2 = \frac{a_2 - a_1}{Aa_1 + 1}U'_2,$$

which has the same sign of U'_2 , as the fraction term is positive by hypothesis.

To complete the proof, we want to show that $U'_2 > 0$. We have the following about unemployment for type 2 workers

$$U_2(\varepsilon) = \frac{\delta N_2}{\delta + \alpha a_2(q(\theta_2(\varepsilon)) + q(\theta_3(\varepsilon)))},$$

$$U'_2 = -\frac{\alpha a_2 U_2 q'(\theta)(\theta'_2 + \theta'_3)}{\delta + 2\alpha a_2 q(\theta)}.$$

Use the values for θ'_2 and θ'_3 from the previous paragraph and re-arrange to obtain

$$U'_2 \left[1 + \underbrace{\frac{\alpha a_2 q'(\theta)}{\delta + 2\alpha a_2 q(\theta)} \left(\theta + \frac{U_2(a_2 - a_1)}{V_3(Aa_1 + 1)} \right)}_{:=B} \right] = \frac{-\alpha a_2 U_2 q'(\theta)\theta}{[\delta + 2\alpha a_2 q(\theta)]V_3}.$$

Notice that the right hand side is positive so $U'_2 > 0$ iff the left hand side is positive too. Below, we are going to show that $B \in (-1, 0)$ which implies the left hand side is positive. The fact that B is negative is trivial as all multiplicative terms are positive except for $q'(\theta)$. To show $B > -1$, observe the following

$$\begin{aligned} B &> \frac{\alpha a_2 q'(\theta)\theta}{\delta + 2\alpha a_2 q(\theta)} \left(1 + \frac{a_2 - a_1}{Aa_1 + 1} \right) \\ &\geq \frac{\alpha a_2 q'(\theta)\theta}{2\alpha a_2 q(\theta)} \left(1 + \frac{a_2 - a_1}{Aa_1 + 1} \right) \\ &> -\frac{q(\theta)}{2q(\theta)} \left(1 + \frac{a_2 - a_1}{Aa_1 + 1} \right) \\ &> -1. \end{aligned}$$

where the first inequality follows from $V_3 > V_2$, the second from $\delta \geq 0$, the third from $q'(\theta)\theta > -q(\theta)$ (as $A > 0$), and the last inequality from Assumption A.1, which guarantees that the term in the big brackets is no larger than 2.

We showed $U'_2 > 0$ which implies E' and O' are positive.

Appendix B: Job Advice Platform




FIGURE B.1. Screen shot of job advice platform - example. Replication of screen shot of advice platform in Belot, Kircher, and Muller (2019) for someone looking for a job in the occupation “waiter/waitress”. The left panel shows the advice from the transition matrix, the right panel the advice based on skill transferability, clicking the heat maps link generates maps of labor demand and as in Figure B.2, and the search button then returns vacancies in all the selected occupations on the screen.




FIGURE B.2. Heat map—example. Replication of screen shot of “heat maps” in Belot, Kircher, and Muller (2019) that can be accessed by job seekers from the general advice represented in Figure B.1. Darker colors indicate a lower ratio of job-seekers to vacancies for a particular set of occupations.

References

- Altmann, Steffen, Armin Falk, Simon Jäger, and Florian Zimmermann (2018). “Learning about Job Search: A Field Experiment with Job Seekers in Germany.” *Journal of Public Economics*, 164, 33–49.
- Banfi, Stefano and Benjamin Villena-Roldan (2019). “Do High-Wage Jobs Attract More Applicants? Directed Search Evidence from the Online Labor Market.” *Journal of Labor Economics*, 37(3), 715–746.
- Behaghel, Luc, Bruno Crépon, and Marc Gurgand (2014). “Private and Public Provision of Counseling to Job Seekers: Evidence from a Large Controlled Experiment.” *American Economic Journal: Applied Economics*, 6, 142–74.
- Belot, Michele, Philipp Kircher, and Paul Muller (2019). “Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice.” *The Review of Economic Studies*, 86, 1411–1447.

- Belot, Michèle, Philipp Kircher, and Paul Muller (2022a). “Do the Long-term Unemployed Benefit from Automated Occupational Advice during Online Job Search?” Center For Economic Policy Research (CEPR) Discussion Paper 17513.
- Belot, Michèle, Philipp Kircher, and Paul Muller (2022b). “Eliciting Time Preferences when Income and Consumption Vary: Theory, Validation and Application to Job Search.” IZA Discussion Paper 14091.
- Belot, Michèle, Philipp Kircher, and Paul Muller (2022c). “How Wage Announcements Affect Job Search Behavior—A Field Experiment.” *American Economic Journal: Macroeconomics*, 14, 1–67.
- Card, David, Jochen Kluve, and Andrea Weber (2018). “What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations.” *Journal of the European Economic Association*, 16, 894–931.
- Cottier, Lionel, Yves Fluckiger, Pierre Kempeneers, and Rafael Lalive (2018). “Does Job Search Assistance Really Raise Employment?” Workin Paper No. 11766, Institute of Labor Economics Available at SSRN 3249880.
- Crépon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora (2013). “Do Labor Market Policies Have Displacement Effects? Evidence from a Clustered Randomized Experiment.” *The Quarterly Journal of Economics*, 128, 531–580.
- Dal Bó, Ernesto, Frederico Finan, and Martín A Rossi (2013). “Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service.” *The Quarterly Journal of Economics*, 128, 1169–1218.
- Dhia, Aïcha Ben, Bruno Crépon, Esther Mbihi, Louise Paul-Delvaux, Bertille Picard, and Vincent Pons (2022). “Can a Website Bring Unemployment Down? Experimental Evidence from France.” National Bureau of Economic Research Working Paper No. 29914.
- Faberman, R. Jason and Marianna Kudlyak (2019). “The Intensity of Job Search and Search Duration.” *American Economic Journal: Macroeconomics*, 11, 327–57.
- Gee, Laura K. (2019). “The More You Know: Information Effects on Job Application Rates in a Large Field Experiment.” *Management Science*, 65, 2077–2094.
- Horton, John J. (2019). “Buyer Uncertainty about Seller Capacity: Causes, Consequences, and a Partial Solution.” *Management Science*, 65, 3518–3540.
- Horton, John J., Ramesh Johari, and Philipp Kircher (2021). “Cheap Talk Messages for Market Design: Theory and Evidence from a Labor Market with Directed Search.” Tech. rep., National Bureau of Economic Research WP 29445.
- Kircher, Philipp A.T. (2020). “Search Design and Online Job Search—New Avenues for Applied and Experimental Research.” *Labour Economics*, 64, 101820.
- Kroft, Kory and Devin G. Pope (2014). “Does Online Search Crowd Out Traditional Search and Improve Matching Efficiency? Evidence from Craigslist.” *Journal of Labor Economics*, 32, 259–303.
- Kuhn, Peter (2014). “The Internet as a Labor Market Matchmaker.” *IZA World of Labor*, Institute of Labor Economics (IZA), 18. doi:10.15185/izawol.18.
- Kuhn, Peter and Kailing Shen (2013). “Gender Discrimination in job ads: Evidence from China.” *The Quarterly Journal of Economics*, 128, 287–336.
- Kuhn, Peter and Kailing Shen (2021). “What Happens When Employers Can No Longer Discriminate in Job Ads?” National Bureau of Economic Research Working Paper No. 29116.
- Kuhn, Peter, Kailing Shen, and Shuo Zhang (2020). “Gender-Targeted Job ads in the Recruitment Process: Facts from a Chinese Job Board.” *Journal of Development Economics*, 147, 102531.
- Marinescu, Ioana and Roland Rathelot (2018). “Mismatch Unemployment and the Geography of Job Search.” *American Economic Journal: Macroeconomics*, 10, 42–70.
- Marinescu, Ioana and Daphné Skandalis (2021). “Unemployment Insurance and Job Search Behavior.” *The Quarterly Journal of Economics*, 136, 887–931.
- Marinescu, Ioana and Ronald Wolthoff (2020). “Opening the Black Box of the Matching Function: The Power of Words.” *Journal of Labor Economics*, 38, 535–568.

- Mueller, Andreas I., Damian Osterwalder, Josef Zweimüller, and Andreas Kettemann (2018). "Vacancy Durations and Entry Wages: Evidence from Linked Vacancy-Employer-Employee Data." National Bureau of Economic Research Working Paper No. 25118.
- Naya, Victor Alfonso, Guillaume Bied, Philippe Caillou, Bruno Crépon, Christophe Gaillac, Elia Pérennes, and Michèle Sebag (2021). "Designing Labor Market Recommender Systems: The Importance of Job Seeker Preferences and Competition." In 4. *IDSC of IZA Workshop: Matching Workers and Jobs Online-New Developments and Opportunities for Social Science and Practice*. Online, France. hal-03540319.