How wage announcements affect job search behaviour - a field experimental investigation

Michele Belot, Philipp Kircher, and Paul Muller*

February 2017

Abstract

In this study we introduce a small number of "artificial" vacancies with randomised characteristics in an otherwise standard job search platform. This allows us to study how job seekers react to job characteristics - in analogy of usual randomised audit studies of employers' reactions to applicant characteristics. We focus on the reaction to wage announcements, and test the main implications of directed search: high wages should attract more and better applicants, but some applicants apply only to low wages even if higher wage offers are present. Both parts of the theory find support among the randomised job offers, suggesting an allocative role for wage competition in search markets. We calibrate a directed search model with multiple applications and on-the-job search and find that it can reproduce our findings quantitatively.

Keywords: Online job search, directed search, wage competition, audit studies, randomised experiments.

JEL-codes: J31, J63, J64, C93

Preliminary and Incomplete
Do not cite/circulate

* *Affiliations*: Belot and Kircher, University of Edinburgh; Muller, University of Gothenburg. We thank the Job Centres' in Edinburgh for their extensive support for this study, and especially Cheryl Kingstree who provided invaluable help and resources. We thank the Applications Division at the University of Edinburgh and in particular Jonathan Mayo for his dedication in programming the job search interface and databases, and to Peter Pratt for his consultation. We thank the UK Minister for Employment Mark Hoban as well as Tony Jolly at the UK Department for Work and Pensions Digital Services Division for granting us access to the vacancy data, and to Christopher Britton at Monster.com for liaising with us to provide the technical access. We are grateful to Andrew Kelloe, Jonny Horn, Robert Richards and Samantha Perussich for extensive research assistance and to Ivan Salter for managing the laboratory. Kircher acknowledges the generous financial support from the European Research Council Grant No. 284119, without which this study would not have been possible.

1 Introduction

Audit studies of employer behaviour have a long tradition in sociology and economics, and have provided important insights in how employers react to applicant characteristics such as skills, age, gender, unemployment duration, and ethnicity. The power of these studies comes from their clean identification: applicants or applications are randomised, so any differences can be attributed to the employer.

This paper advocates a similar approach to study the reaction of job seekers to differences in the vacancies that are posted. By placing a small set of research vacancies with randomised characteristics such as the wage offer, one can study how wage competition works, and test different theories of the labor market. In particular, we aim to provide clean evidence on two assumptions that set models of direct wage competition (competitive or directed search) apart from models where workers encounter wages more randomly (random search): (1) job seekers do not have the time to apply to all job offers and sometimes apply to low wages even if high wage offers are present, but (2) overall more and/or better workers apply to jobs with higher wages. The second point implies that wages have a competitive character in attracting workers, which generates positive welfare theorems akin to classical general equilibrium theory in competitive search models.[1] This has been the main reason why these theories have attracted larger interest as an alternative to random search. The first point is also crucial for the theory: if workers send applications to all jobs that surpass their reservation wage the main selectivity argument underlying these theories is lost.[2]

Of course it is empirically challenging to assess the effects of wages on applications, and this is why we propose an experimental approach. In order to cleanly identify the effects of wages, one needs to ensure that one is comparing jobs that are otherwise comparable. This is usually difficult to do with existing data sets. Our experimental approach circumvents some of these challenges.

The general methodology we propose could also be used to study wider questions, such as whether minorities react differently to higher wages or employer ethnicity or gender. A limited number of research vacancies would allow to shed light on this additional source of minority wage differentials if these could reach a sufficiently large and diverse set of job seekers. Unfortunately current ethical approval practices did not allow this approach to proceed without prior informed consent, even after laying out procedures that substantially curb the downsides.[3] Consent of all market participants is unfeasible, similar to standard audit studies of employers. In this paper we aim to overcome this hurdle in a small scale field experiment to demonstrate the benefits of such an experimental approach

[1] Related literature.

[2] An example where higher wages attract more applicants but search is undirected is the following slight variation on undirected on-the-job search models: Assume that workers encounter wage offers by firms randomly as in Burdett and Mortensen (1998) but workers only bother to send a formal application if the wage is above their reservation wage. All outcomes are identical to those in Burdett and Mortensen (1998), but higher wage firms attract more applicants if applicants differ in their utility of leisure.

[3] Research vacancies were to be posted only in occupations with many job seekers and many applicants so that rejections to any particular vacancy are the norm rather than a source of major frustration. Survey evidence indicates that most individuals only search a few minutes per day, and by focusing on occupations where this is in the lower end avoids high costs of wasted effort. Many online platforms feature a non-trivial fraction of fictitious jobs, which websites for job search advice openly discuss. Relative to these our intervention would have been small in scale. We elaborate on these considerations in the conclusions.

to job seeker behaviour. We focus on the general questions in the preceding paragraph, since the restrictions limit our power to answer the more detailed questions outlined here which will hopefully become feasible in future studies.

We created our own job matching platform at the experimental laboratory in Edinburgh. We recruited 300 job seekers to participate in our study once a week for a duration of 12 weeks. Per session they spend at least half an hour searching for jobs that they can save to apply later, and they can spend up to two hours to actually apply or they can do the actual applications from home. Conducting the study in our lab allows us to ask participants explicitly for consent, to verify their identities, to distribute compensation, and to provide support in case of platform problems. For the other side of the market we provided access to some 600.000 up-to-date vacancies from the database of Universal Jobmatch, the UK government's job search site which is the most comprehensive in the UK market. We informed participants that a small fraction (less than 2% of vacancies) would be posted for research purposes to understand which vacancies should be attracted to the region.[4] Job seekers were informed of the source of our regular vacancies. We did not specify the exact nature of our research vacancies.

In praxis, these were expired vacancies from half a year prior to our study that did not allow direct name identification and for which we artificially multiplied the original wage by some random factor. Since the multiplicative factor is assigned independent of the description of the vacancy, we can answer the second question by regressing the amount of interest and the qualification of the interested parties on the wage factor. We created pairs of vacancies that are identical except for minor changes in wording and are posted at roughly the same time but with different wages. This allows us to answer the first research question whether there exist people who apply only to the low wage even if higher wage jobs in the same area are present (even though the higher wage should on average attract more/better applicants). The random assignment of wages is crucial to our identification strategy, as wages are usually correlated with other indicators of hierarchy or status that are conveyed in the flow text and that are difficult to control for.[5]

We do not expect or have evidence that the presence of some non-real vacancies alters the job search of our participants substantially. Even in the absence of our postings, online datasets are not pristine collections of open vacancies. Newspaper reports of and investigations into fake job advertisements about websites like Monster.com, Careerbuilder.com and Universal Jobmatch, and a large number of online advice websites inform job seekers of this general problem. For our particular vacancy source form Universal Jobmatch, investigative broadsheet journalism in the UK places the lower bound of non-real vacancies at 2%, while the tabloid press warns of much higher numbers.[6] Nevertheless, the database is regarded as reliable by the UK government, who considers making search on this platform mandatory for Job Seekers Allowance recipients in the future. We take from this that a low number of non-real vacancies are standard in current online job search. We also survey our participants at the end of the study, and the vast majority specified that the presence of research vacancies was immaterial for how they search for jobs, nor

[4] We chose this formulation based on the interest of policy makers in understanding how important more attractive vacancies are for job seekers relative to providing more jobs at all even at lower wages.

[5] For example, even in the study by Marinescu and Wolthoff (2014), which includes job titles fixed effects, two lines of flow text remain which cannot be controlled for.

[6] See for example Channel 4 (2014) and Computer Business Review (2014)

do they deem it possible to distinguish our vacancies from real ones.

As measure of interest we record both whether individuals view the details of a particular vacancy and whether they save it. After saving a vacancy they can apply for it once they leave the search phase after a minimum of half an hour. For ethical reasons we informed participants right after the initial search phase, and therefore prior to actual applications, if they saved any vacancies that we had posted for research purposes. This avoids unnecessary work and anxiety of applying to jobs that are not real. It also ensures that they never see any research vacancy twice, as we delete it from their search set thereafter. It implies that we do not observe actual applications, however viewing or saving a vacancy is a strong indicator of interest in the vacancy and likelihood of applying.[7]

Our main results are in line with the two main predictions of directed search models. First, we find that higher wages result in significantly more interest in the vacancy. A 1% increase in the wage results in 0.34% more viewings. Second, we find that 37% of those that view the low wage vacancy within a vacancy-pair do not also view the high wage vacancy. This percentage remains almost constant when conditioning on individuals that have seen both vacancies on their screen (36%). It becomes somewhat smaller when also conditioning on pairs in which the high wage vacancy was posted later (and thus appeared higher in the list), though even in this scenario 20% does not view the high wage vacancy. In a robustness analysis we exclude that this finding is due to study participants identifying the research vacancies or because of locational differences between the vacancies in a pair.

Using a complementary survey (with different participants), we obtain perceptions about the artificial vacancies. We find that the high wage vacancy within a pair is, on average, perceived to (1) attract more competition, (2) require an applicant to be of higher quality to be considered and (3) have better non-monetary working conditions. Findings (1) and (2) support a directed search interpretation of our empirical results, while finding (3) rules out a compensating differential hypothesis (in which workers do not apply to high wage jobs because they expect non-monetary conditions to be worse).

To provide a benchmark for the magnitude of our estimates, we build a directed search model with multiple applications and on-the-job search. We calibrate the model using UK data and compute the variables of interest: the queue length elasticity of a vacancy with respect to its posted wage and the worker's probability of applying to low-wage job but not to a high wage job when both are observed. We show that the model is able to reproduce values reasonably close to our empirical findings.

The rest of the paper is organized as follows. Section 2 provides an overview of related literature. In Section 3 we describe the set up of the experiment and present some descriptive statistics on the participating job seekers. The details on how the artificial vacancies were created and posted are explained in section 4. Section 5 presents the empirical analysis and results. In Section 6 a simple directed search model is described and its predictions are compared to our empirical findings. Section 7 concludes.

2 Related Literature

[7] In our study, almost one-third of all saved (real) vacancies is eventually applied to.

Our study considers how the posted wage of a vacancy affects the decision of job seekers on whether or not to consider it for application. This relationship is generally considered to be confounded by many factors. Naturally, one has to control for a set of vacancy characteristics such as occupation and qualification requirements. But even conditional on a large set of vacancy characteristics, vacancies with different wages might differ in other respect that might be conveyed through the flow text or the specific job title. Holzer et al. (1991) analyze applications data and find that firms that are restricted to pay a minimum wage receive more applicants that firms that pay either slightly less or slightly more than the minimum wage. Marinescu and Wolthoff (2014) use online job search data from Careerbuilder.com and show that higher wages are associated with fewer applicants, even when controlling for occupation and industry. Only when conditioning on the exact job title, the association reverses and higher wage atrract more applicants. Also the quality of the applicants (education and experience) increases with the posted wage. In order to overcome endogeneity issues in observational data, a study by Dal Bó et al. (2013) randomizes the wage offers in job openings for civil servants in Mexico. They find that higher wages lead to significantly more applicants, while higher wages also increase the quality of the applicants as measured by IQ level.

We propose to apply a methodology that is parallel to audit studies in which the success of pairs of applicants is compared, when they differ in only one particular characteristic. Originally, studies were performed in which applicant pairs were matched to be as closely comparable when applying for a job opening (see for example Neumark et al. (1996), who consider gender discrimination).

These studies were criticized for a number of reasons, leading to a new approach in which résumes where send out. The famous study by Bertrand and Mullainathan (2004) considers discrimination in the hiring process, by varying the names of the applicants between 'African-American'- and 'White-American'-sounding. This approach has become popular and has been applied to identify discrimination based on different grounds such as ethnicity, gender or sexual orientation.

We apply a similar approach to the other side of the labor market, by posting pairs of vacancies that are the same in all respects except for the wage. As such we solve the problem of endogeneity of wages when testing whether wages affect applicant behavior. Even more, the experimental set up allows to test further implications of the directed search model (Moen (1997), Burdett and Mortensen (1998)) as we can observe in detail the job search behavior of individual job seekers.

3 Experimental Design

3.1 Recruitment Procedure and Experimental Sample

The participants in the study were job seekers recruited in the area of Edinburgh. The eligibility criteria were: being unemployed, searching for a job for less than 12 weeks (a criterion that we did not enforce), and being above 18 years old. We imposed no further restrictions in terms of nationality, gender, age or ethnicity. Most participants received unemployment benefits (Job Seekers Allowance). Details about these benefits and corresponding job search requirements can be found in our companion paper (Belot

Table 1: Recruitment and show-up of participants

	Full sample	Wave 1	Wave2
Recruitment channel participants:			
Job centres	86%	83%	89%
Gumtree or other	14%	17%	11%
Sign up rate jobcentre[a]	43%	39%	47%[b]
Show up rate	45%	43%	46%

[a] Of those people that were willing to talk to us about the study, this is the share that signed up for the study. [b] Based on only one day of recruitment (see footnote 9).

et al. (2015)).

We obtained the collaboration of several local public unemployment agencies (Job Centres) to recruit job seekers on their premises. We also placed posters and advertisements at various public places in Edinburgh (including libraries and cafes) and posted a classified ad on a popular on-line platform (not limited to job advertisements) called Gumtree. In Table 1 sign up and show up rates are presented. Of all participants, 86% were recruited in the Jobcentres. Most of the other participants were recruited through our ad on Gumtree. We approached all visitors at the Jobcentres during two weeks.[8] Out of those we could talk to and who did not indicate ineligibility, 43% percent signed up. Out of everyone that signed up, 45% showed up in the first week and participated in the study. These figures display no statistically significant difference between the two waves of the study. In Figure 9 in the Appendix, we show aggregate labor market statistics, which indicate that the participants in the study constitute about 2% of the stock of unemployed in Edinburgh.[9]

The baseline survey at the first week of the study provides a number of background characteristics for the participants, which can be compared to statistics on the population of job seekers in Edinburgh. Both are presented in Table 2. Unfortunately only a limited number of characteristics is available for the population of job seekers in Edinburgh, which are based on the NOMIS dataset on Job Seeker Allowance (JSA) Claimants. We find that we slightly oversample females and non-whites, while the average age is very close the population average. The population statistics are based on JSA claimants with unemployment duration up to 6 months, because for these the median duration is almost equal to the median duration of the participants (80 days). For the participants we further observe that they have, on average send out 64 applications and attended 0.48

[8] Since most Job Seekers Allowance recipients were required to meet with a case worker once every two weeks at the Jobcentre, we were able to approach a large share of all job seekers.

[9] We asked the recruiters to write down the number of people they talked to and the number that signed up. Unfortunately these have not been separated for the online study and the lab study. In the first wave there were different recruiters for the two studies, such that we can compute the sign up shares separately. In the second wave we asked assistants to spend parts of their time per day exclusively on the lab study and parts exclusively on the online study, so we only have sign-ups for the total number. One day was an exception, as recruitment was done only for the lab study on this day, such that we can report a separate percentage based on this day.

job interviews.

Table 2: Characteristics of lab participants

	Lab participants				Population[a]
	mean	sd	min	max	
Demographics:					
gender (%)	43	50	0	1	33
age	36	12	18	64	35
high educated[b] (%)	43	50	0	1	
white (%)	80	40	0	1	89
number of children	.53	1	0	5	
couple (%)	23	42	0	1	
any children (%)	27	45	0	1	
Job search history:					
vacancies applied for	64	140	0	1000	
interviews attended	.48	0.84	0	6	
jobs offered	.42	1.1	0	8	
at least one offer (%)	20	40	0	1	
days unempl. (mean)	260	620	1	5141	111
days unempl. (median)	80				81
less than 183 days (%)	76	43	0	1	
less than 366 days (%)	85	35	0	1	
job seekers allowance (£)	52	75	0	1005	
housing benefits (£)	64	129	0	660	
other benefits (£)	14	65	0	700	
Observations	295				

[a] Average characteristics of the population of job seeker allowance claimants in Edinburgh over the 6 months of study. The numbers are based on NOMIS statistics, conditional on unemployment duration up to one year. [b] High educated is defined as a university degree.

3.2 Experimental Procedure

Job seekers were invited to search for jobs once a week for a period of 12 weeks (or until they found a job) in the computer facilities of the School of Economics at the University of Edinburgh. The study consisted of two waves: wave 1 started in September 2013 and wave 2 started in January 2014. We conducted sessions at six different time slots, on Mondays or Tuesdays at 10 am, 1 pm or 3:30 pm. Participants chose a slot at the time of recruitment and were asked to keep the same time slot for the twelve consecutive weeks.

Participants were asked to search for jobs using our job search engine (described later in this section) for a minimum of 30 minutes.[10] After this period they could continue to search or use the computers for other purposes such as writing emails, updating their CV,

[10]The 30 minute minimum was chosen as a trade-off between on the one hand minimizing the effect

or applying for jobs. They could stay in our facility for up to two hours. All participants received a compensation of £11 per session attended (corresponding to the government authorized compensation for meal and travel expenses) and we provided an additional £50 clothing voucher for job market attire for participating in 4 sessions in a row.[11]

Participants were asked to register in a dedicated office at the beginning of each session. At the first session, they received a unique username and password and were told to sit at one of the computer desks in the computer laboratory. The computer laboratory was the experimental laboratory located at the School of Economics at the University of Edinburgh with panels separating desks to minimize interactions between job seekers. They received a document describing the study as well as a consent form that we collected before the start of the initial session (the form can be found in the Online Appendix OA.1). We handed out instructions on how to use the interface, which we also read aloud (the instructions can be found in the Online Appendix OA.2). We had assistance in the laboratory to answer questions. We clarified that we were unable to provide any specific help for their job search, and explicitly asked them to search as they normally would.

Once they logged in, they were automatically directed to our own website.[12] They were first asked to fill in a survey. The initial survey asked about basic demographics, employment and unemployment histories as well as beliefs and perceptions about employment prospects. From week 2 onwards, they only had to complete a short weekly survey asking about job search activities and outcomes. For vacancies saved in their search in our facility we asked about the status (applied, interviewed, job offered). We asked similar questions about their search through other channels than our study.

After completing the survey, the participants were re-directed towards our search engine and could start searching. A timer located on top of the screen indicated how much time they had been searching. Once the 30 minutes were over, they could end the session. They would then see a list of all the vacancies they had saved and were offered the option of printing these saved vacancies. We received no additional information about the search activities or search outcomes from the Jobcentres. We only received information from the job seekers themselves. This absence of linkage was important to ensure that job seekers did not feel that their search activity in our laboratory was monitored by the employment agency. They could then leave the facilities and receive their weekly compensation.[13] Those who stayed could either keep searching with our job search engine or use the computer for other purposes (such as updating their CV, applying on-line or using other job search engines). Once participants left the facility, they could still access our website from home, for example in order to apply for the jobs they had found.[14]

---

of participation on the natural amount of job search, while on the other hand ensuring that we obtained enough information. Given that participants spent around 12 hours a week on job search, a minimum of half an hour per week is unlikely to be a binding constraint on weekly job search, while it was a sufficient duration for us to collect data.

[11] All forms of compensation effectively consisted of subsidies, i.e. they had no effect on the allowances the job seekers were entitled to. The nature and level of the compensation were discussed with the local job centres to be in accordance with the UK regulations of job seeker allowances.

[12] www.jobsearchstudy.ed.ac.uk

[13] Participants were of course allowed to leave at any point in time but they were only eligible to receive the weekly compensation if they had spent 30 minutes searching for jobs using our search engine.

[14] When searching from home participants would never encounter artificial vacancies in their search results. Artificial vacancies were only included in the set of vacancies during the experiment sessions.

Figure 1: Number of vacancies

(a) Posted vacancies in our study

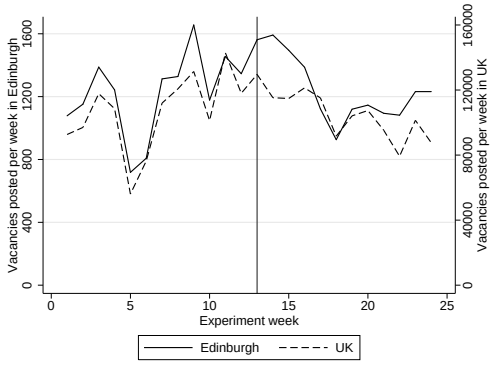

(b) Active vacancies in our study and in UK

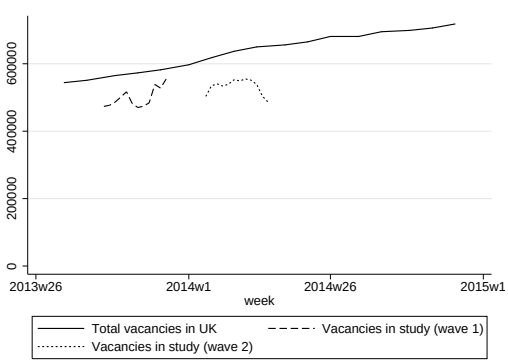

3.3 The search interface

We designed a job search engine in collaboration with the computer applications team at the University of Edinburgh. It was designed to replicate the search options available at the most popular search engines in the UK (such as monster.com and Universal Jobmatch), but allowing us to record precise information about how people search for jobs (what criteria they use, how many searches they perform, what vacancies they click on and what vacancies they save), as well as collecting weekly information (via the weekly survey) about outcomes of applications and search activities outside the laboratory.

In order to provide a realistic job search environment, the search engine accesses a local copy of the database of real job vacancies of the government website Universal Jobmatch. This is the largest job search website in the UK in terms of the number of vacancies. This is a crucial aspect in the setup of the study, because results can only be trusted to resemble natural job search if participants use the lab sessions for their actual job search. The large set of available vacancies combined with our carefully designed job search engine assures that the setting was as realistic as possible. Panel (a) of Figure 1 shows the number of posted vacancies available through our search engine in Edinburgh and in the UK for each week of the study (the vertical line indicates the start of wave 2). Each week there are between 800 and 1600 new vacancies posted in the Edinburgh. Furthermore, there is strong correlation between vacancy posting in Edinburgh and the UK. In panel (b) the total number of active vacancies in the UK is shown over the second half of 2013 and 2014.[15] As a comparison the total number of active vacancies in the database used in the study in both waves is shown. It suggests that the database contains over 80% of all UK vacancies, which is a very extensive coverage compared to other online platforms.[16] It is well-known that not all vacancies on online job search platforms are

[15] Panel (b) is based on data from our study and data from the Vacancy Survey of the Office of National Statistics (ONS), dataset "Claimant Count and Vacancies - Vacancies", url: www.ons.gov.uk/ons/rel/lms/labour-market-statistics/march-2015/table-vacs01.xls

[16] For comparison, the largest US jobsearch platform has 35% of the official vacancies; see Marinescu (2014), Marinescu and Wolthoff (2014) and Marinescu and Rathelot (2014). The size difference might be due to the fact that the UK platform is run by the UK government.

Figure 2: Standard search interface

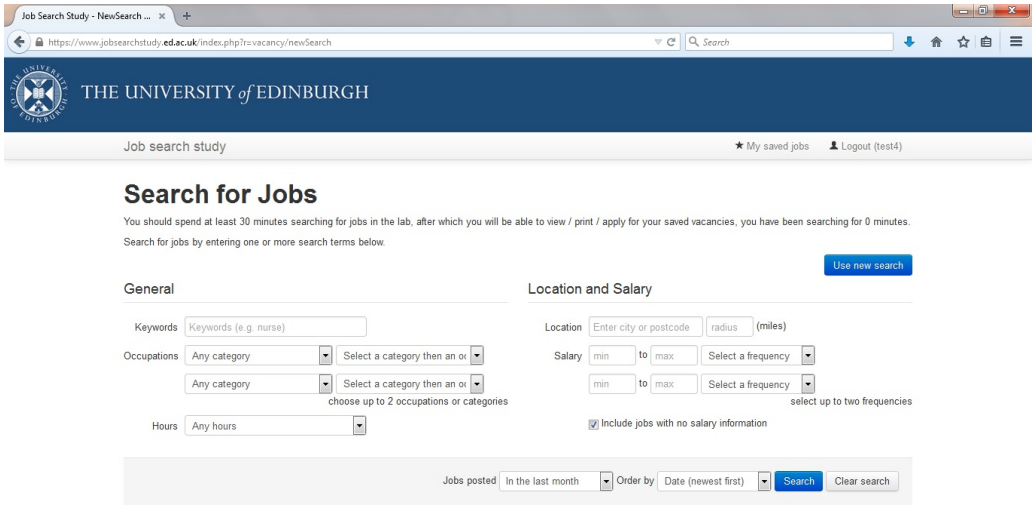

genuine, so the actual number might be somewhat lower.[17]

Figure 2 shows a screenshot of the main page of the standard search interface. Participants can search using various criteria (keywords, occupations, location, salary, preferred hours), but do not have to specify all of these. Once they have defined their search criteria, they can press the search button at the bottom of the screen and a list of vacancies fitting their criteria will appear. The information appearing on the listing is the posting date, the title of the job, the company name, the salary (if specified) and the location. They can then click on each individual vacancy to reveal more information. Next, they can either choose to "save the job" (if interested in applying) or "do not save the job" (if not interested). If they choose not to save the job, they are asked to indicate why they are not interested in the job from a list of possible answers.

As in most job search engines, they can modify their search criteria at any point and launch a new search. Participants had access to their profile and saved vacancies at any point in time outside the laboratory, using their login details.

From week 4 onward, half of the participants were offered to use an "alternative" interface which was designed to investigate how occupational broadness in job search affects job prospects. Since it is not directly related to the research question addressed in this paper, we only briefly describe the "alternative" interface here. An extensive description as well as an empirical analysis of the impact of the interface can be found in Belot et al. (2015). The key goal of the alternative interface was to offer suggestions to job seekers about occupations that might be of interest to them. This was achieved by

[17] For Universal Jobmatch evidence has been reported on fake vacancies covering 2% of the stock posted by a single account (Channel 4 (2014)) and speculations of higher total numbers of fake jobs circulate (Computer Business Review (2014)). Fishing for CV's and potential scams are common on many sites, including Carreerbuilder.com (The New York Times (2009a)) and Craigslist, whose chief executive, Jim Buckmaster, is reported to say that "it is virtually impossible to keep every scam from traversing an Internet site that 50 million people are using each month" (The New York Times (2009b)).

creating a list of potentially interesting occupations, based on the preferred occupation of the participant. Two methodologies were applied to create this list. First, labor market surveys were used to identify the most common transitions between occupations. Second, occupations that require the same set of skills as the preferred occupation (based on the US based website O*Net) were suggested. Participants select which suggestions they find interesting after which a search is performed for all selected occupations.

Even though the alternative interface affects individual job search behavior, it is orthogonal to the randomized set up of the artificial vacancies. Therefore it does not interfere with the empirical analysis. We discuss this issue in more detail in section 5 and present robustness checks where only participants that used the standard interface are included.

4 Artificial vacancies

A small number of artificial vacancies was introduced during the study. Participants were fully informed about this. They were told that "we introduced a number of vacancies (about 2% of the database) for research purposes to learn whether they would find these vacancies attractive and would consider applying to them if they were available". Participants were asked for consent to this small percentage of research vacancies at the start of the study.[18] The artificial vacancies were created and posted on a weekly basis, where the number was determined such that the overall share of artificial vacancies in the stock of vacancy in the Edinburgh area never exceeded 2%. We also checked whether the share of artificial vacancies within all vacancies saved by participants did not exceed 2%, and adjusted the number in subsequent weeks in case it did. The vacancies were added to the database of real vacancies during the days on which lab sessions for participants took place, however in case participants searched from home they could not encounter them. Each artificial vacancy was only active during one particular week, such that participants would never observe them in multiple sessions. In this section we describe the procedure used to create the artificial vacancies and present some statistics on how comparable they are to the set of real vacancies.

4.1 Selection procedure and representativeness

The artificial vacancies were produced in the following manner. We selected an old set of real vacancies that were posted in the UK on Universal Jobmatch during the summer of 2013, which is several months before our study started. From these we selected all vacancies with a wage indication (either a minimum or a maximum wage or both). No restriction was made on whether these were hourly, weekly, monthly or annual salary indications. From this set of vacancies we selected vacancies to use as templates for the artificial vacancies. One key restriction in this process was that the vacancy would have to be sufficiently compact and general in order to be easily manipulated and remain unidentifiable. This restriction is likely to lead to a selective bias towards lower-skilled vacancies (with less extensive vacancy text etc.). From each selected vacancy we removed

[18] In an exit survey the vast majority of participants (86%) said that this did not affect their search behavior. This is likely due to the very low numbers of artificial vacancies and to the fact that fake advertisements are common in any case to online job search sites (see footnote 17) and that this is mentioned to job seekers in many search guidelines (see e.g. Joyce (2015)).

Figure 3: Occupational distribution of vacancies

(a) Real vacancies

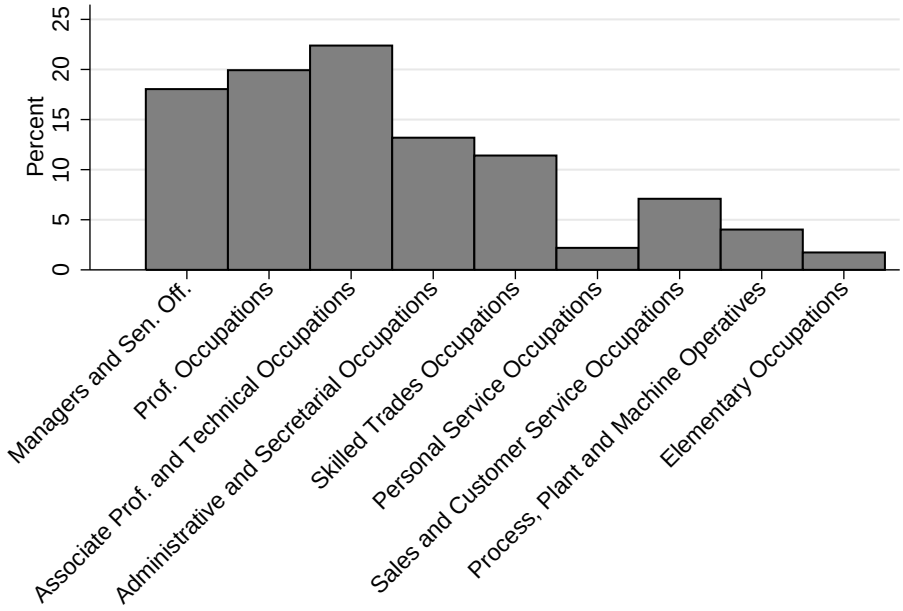

(b) Real vacancies with posted salary

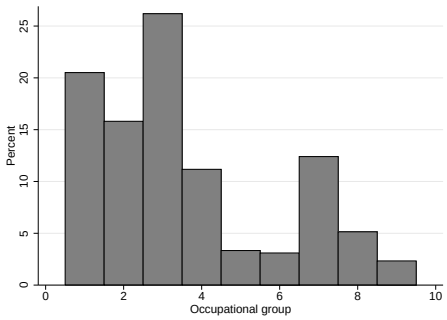

(c) Artificial vacancies

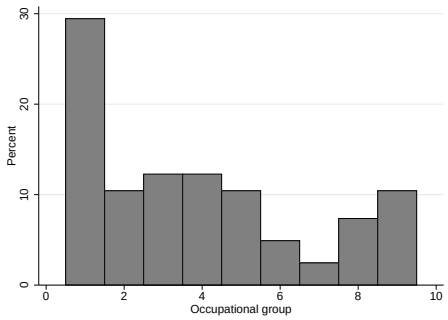

Figure 4: Posted salary: annual

(a) Real vacancies

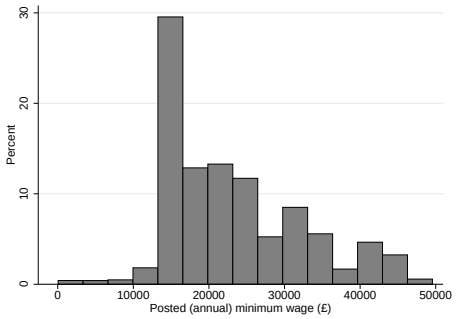

(b) Artificial vacancies

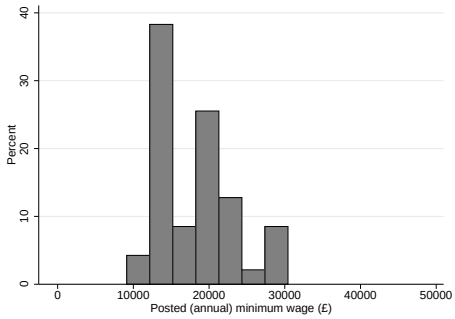

all identifying information (company name, contact person, telephone number, website, etc.).[19] Subsequently we randomly changed the location and the salary of the vacancy, the details of this step are described in the next section. First we discuss to what extent the artificial vacancies are representative of real vacancies.

Given the selection procedure for creating the artificial vacancies, these are likely to differ somewhat from the distribution of real vacancies. In order to manipulate the salary, we required the vacancy to post some salary. Approximately 58% of all vacancies on Universal Jobmatch post a salary, and vacancies that post salaries may differ from those that do not.[20] Panel (a) of Figure 3 shows the distribution across occupations, of vacancies that were posted on Universal Jobmatch during the study.[21] The vacancies are classified by the first digit of their UK SOC code. We present the same distribution for the selection of vacancies that post a salary in panel (b) of Figure 3 .The distribution of vacancies with posted salaries is quite similar to the overall distribution, with only occupational groups 3 and 7 being more likely to post a salary. The second step in the selection procedure required vacancies to have a 'simple' description that allows easy manipulation. The artificial vacancies that we created fulfill this criteria, and their occupational distribution is presented in panel (c) of Figure 3). They differ in two ways. We oversample vacancies from occupational group 1 (Managers and senior officials) and from groups 8 and 9 (both lower skilled jobs). The latter is likely to result from the fact that these are typically shorter and less identifiable job descriptions, making them likely to be suitable for manipulation. Still, there is considerable variation across occupations in the artificial vacancies.

Employers can post an hourly, daily, weekly, monthly or annual wage in their vacancy. Since hourly and annual salaries are most common, we show the distribution of salaries for these two types in Figures 4 and 5. In both figures panel (a) shows the distribution

[19] We also made sure that applying to the vacancy would go through an integrated button saying 'apply now' (which is quite common on Universal Jobmatch) rather than by directly contacting the company or through a company website.

[20] Note that this share is considerably higher than for example documented by Marinescu and Wolthoff (2014), who find that only 20% of vacancies on careerbuilder.com specify wages.

[21] It is based on a sample of 30,000 vacancies posted in the Edinburgh area around the start of the study.

Figure 5: Posted salary: hourly

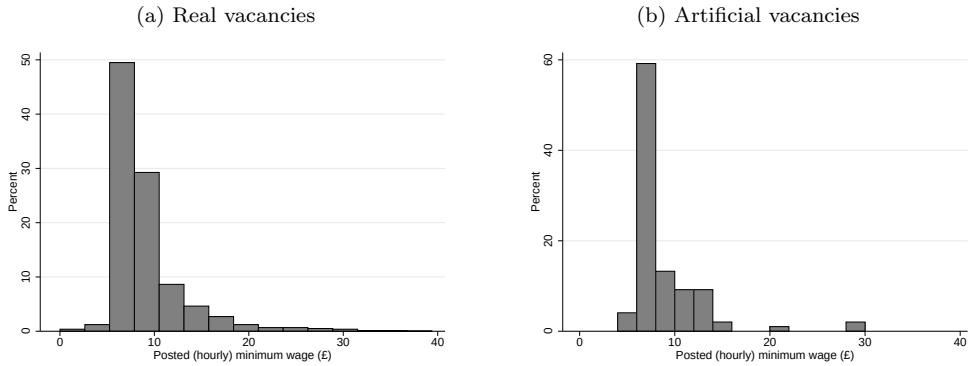

of real vacancies, while panel (b) shows the distribution of the artificial vacancies (before manipulating the wage). From both comparisons it is clear that the artificial vacancies lack some mass in the tails of the distribution, but other than that they are quite similar.[22]

4.2 Manipulation of wages and locations

The goal of producing artificial vacancies is to assess how job seekers respond to the posted wage in a vacancy. It is typically difficult to estimate the effect of wage postings on the number of applicants as the posted wage is correlated with many other attributes of a vacancy. This correlation remains even when controlling for many vacancy attributes, as there are often 'unobserved' factors such as the wording of the flow text or the job title which convey additional information to the job seeker (Marinescu and Wolthoff (2014)).

Our strategy is to create variation in the posted salary while keeping all other vacancy characteristics constant. This could in theory be achieved by posting artificial vacancies with salaries that randomly deviate from the original wage. Such a strategy would however require an enormous sample size to identify the effect of the wage differential on the number of applicants. Alternatively, we create pairs of vacancies. Both share all key vacancy attributes, and differ only in terms of the posted vacancy.[23]

This approach is parallel to the randomized audit studies in which pairs of applicant's résumes are sent out with random variation in one particular dimension. To be able to test the implications of directed search directly, we decided to make both vacancies accessible to the same job seeker. This is in contrast to the resume audit studies, where typically employers are only sent one of the résumes from a pair. The other resume is sent to a different employer. We make both vacancies accessible here because that allows us to observe whether job seekers considers both vacancies or only one of them and in case

[22] Only vacancies with annual salaries up to £50000 or hourly wages up to £40 are shown in the Figure. This excludes 7.7% (annual wage) and 3.4% (hourly wage) of the vacancies.

[23] Our approach is close to 'cluster-randomization', which proposes to match units based on observed characteristics and randomly assign treatment within the matched pair. This allows for efficiency gains compared to unconditional randomization (see for example Imai et al. (2009)). Our approach improves on standard applications of cluster randomization, as we are able to create two identical units that only differ in treatment status.

Figure 6: Salary differences of artificial vacancies

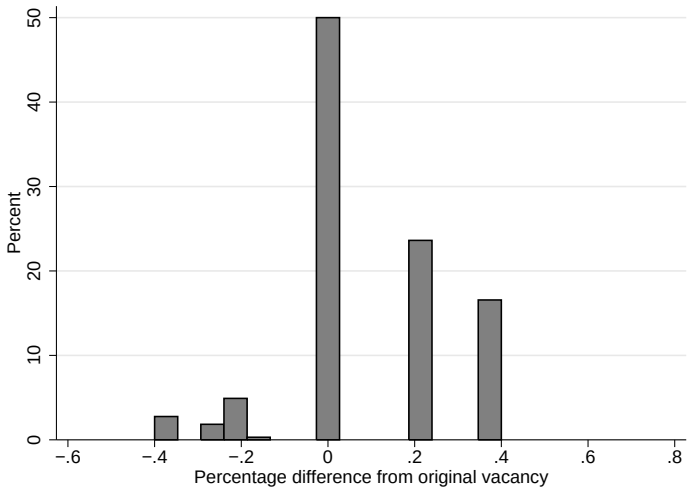

of the latter, which one of the two (s)he chooses to leave out. The use of pairs of artificial vacancies allows to compare the number of applicants within the pair, thereby controlling for all unobserved characteristics of the vacancy. This greatly improves precision of such an estimate. We rephrase and shuffle around the descriptive text of the vacancies in a pair to make sure it is not obvious that they are the same. The key point is that the information conveyed by the two vacancies is the same and the change in the posted salary is independent of other vacancy characteristics. Our setup even allows to check whether a particular job seeker listed both vacancies and thus made a conscious decision to apply to only one of the two.

The construction of the artificial vacancy pairs was done in the following manner. We created pairs of vacancies from the same template vacancy, and for both vacancies we changed the location to the Edinburgh area (with a random postal code). One of the two would keep the original salary, the other one would have a lower or higher wage (20% or 40%). There were two stages of randomization here. First it was randomly decided which vacancy would have a changed salary, second it was randomly decided what the salary change would be. We made sure however, that in case of a salary reduction the new salary would not be below the minimum wage.[24] As result our sample contains relatively more wage increases among low wage vacancies than among high wage vacancies, while also overall we have more wage increases than decreases.[25] During the second wave of

[24] We did so to prevent the vacancy from looking suspicious, though the set of real vacancies actually contains posted wages below the minimum wage as well.

[25] We also created pairs similar to the ones described, with location being the Glasgow area, which is located at about 1.5 hours of commuting time from Edinburgh. Since the willingness of our participants to apply to jobs in the Glasgow area is very small, we have few observations for the these pairs and we focus our analysis on the Edinburgh pairs. Third, we also created some "single" vacancies located in either Edinburgh or Glasgow. These were merely created to make sure participants would not be able to "detect" artificial vacancies from the fact that there were two somewhat similar vacancies. Also for these vacancies we randomly changed the salary. Finally, we created pairs of vacancies, where one would

Figure 7: Salary changes in terms of standard deviations across 3-digit occupations

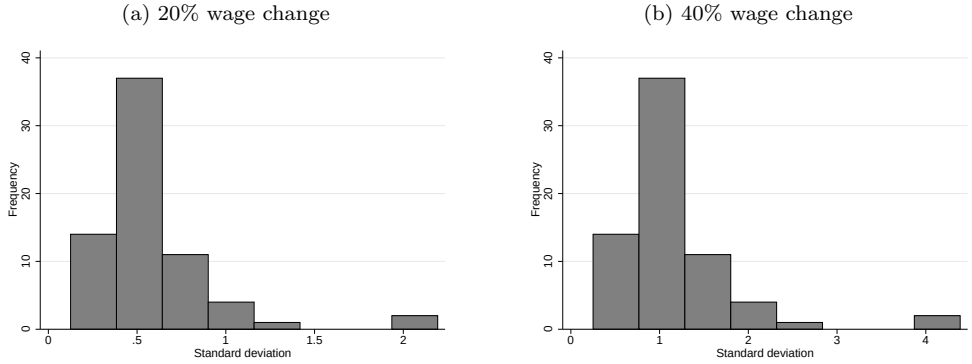

the study, the same set of artificial vacancies was used, however the wage was switched around within the pair.

We show the distribution of the difference in the posted salary compared to the salary of the *original* vacancy if Figure 6. Panel (a) contains all "Edinburgh pair" artificial vacancies. Since one vacancy in each pair has the original salary, 50 % has a wage differential of 0. The manipulated wages are spread across increases and decreases of 20% or 40%, with more increases than decreases.[26]

It is important that the artificial vacancies do not appear suspicious to potential applicants. Therefore the manipulated wages should be reasonable in the sense that they should be within the support of the distribution of wages for the particular occupation. To show that this is the case, we present a measure of wage dispersion at the 3-digit occupational code level. We compute for each 3 digit code the magnitude of a 20 or 40% wage change in terms of the standard deviation of wages.[27] We find that on average a 20% wage increase or decrease corresponds to 0.44 of a standard deviation, while a 40% wage changes corresponds to 0.88 of a standard deviation. The distribution of these numbers across occupations is shown in Figure 7. For almost all occupations a 20 or 40% change in the wage is not likely to be outside the support of the wage distribution.

At the end of the study we performed a small survey to assess whether participants felt that the artificial vacancies had affected their behavior. When asked whether they were able to distinguish the artificial vacancies from real vacancies, 68% answered 'never' or 'rarely'. Only 4% said they could 'often' distinguish them, while 26% answered 'sometimes'. We also asked whether the existence of the artificial vacancies changed their job search behavior. 86% said that it had no effect, 11% answered to save somewhat more vacancies and 1% saved less vacancies. So overall, it did not seem to have a large effect on job search behavior. In section 5.3.4 we show that there is indication of any learning

be located in the Edinburgh area and one in the Glasgow area. For these pairs, either the salary of the vacancy in Glasgow would be increased by 10, 20 or 30%. In the appendix we show that including the Glasgow pairs does not change our results.

[26] Initially we only created artificial vacancies with increased wages, while later we decided to also add some decreased wage vacancies. In addition, not all wage decreases were feasible due to the minimum wage lower bound. As a result we have many more wage increases than decreases.

[27] These computations are based on vacancies that post a minimum annual wage above 1000£.

among job seekers in terms of identifying artificial vacancies.

4.3 Outcome variables

The search process was structured as follows. After specifying search criteria the job seeker observed a list of search results ("listed vacancies"). If a particular vacancy seemed interesting (s)he could click on the vacancy to view the detailed description of the vacancy ("viewed vacancies"). After reading the details, (s)he could save the vacancy to apply later ("saved vacancies"). In case the vacancy was not interesting (s)he could return to the search results after indicating from a list of options why the vacancy was not interesting. At the end of the session they receive the list of saved vacancies. In case their list contains artificial vacancies they would, at this point, be informed about the nature of these vacancies.

Our main analysis focuses on the decision to view of vacancy. This is a clear signal of interest and occurs after observing the key characteristics of the vacancy (job title, location and salary). Clearly, the saving decision is of interest as well, and can be argued to be a closer proxy to the decision to apply. However, once the vacancy is viewed, job seekers may be able to identify pairs of vacancies due to their great similarity. As a result, the decision of whether or not to save the vacancy might be affected by whether or not the job seeker suspects the vacancy to be the same as a previous one.[28] This concern is absent for "viewed vacancies" as the combination of a job title, location and salary are unlikely to identify one particular job opening. This said, we will also present results on "saved vacancies".[29]

So our measure of interest is the number of times a vacancy was viewed. The distribution of the number of times an artificial vacancy was viewed is shown in Figure 8.[30] 22% is never viewed (73 vacancies), 50% is viewed between 1 and 5 times (163 vacancies), and 28% is viewed more than 5 times (90 vacancies). The mean number of views for Edinburgh vacancies is 3.6. Summary statistics for the number of saved vacancies are presented in the Appendix (Figure **??**).

5 Empirical analysis

We now turn to the empirical analysis. We will focus on our two main research questions: First, do we find evidence that higher wages increase interest in vacancies, all else equal? Second, do we find evidence for the reservation wage property, that is an inherent property of random search models but is violated in directed search models?

[28] We find some evidence that this is the case from the responses after not saving a viewed vacancy. A significantly larger share indicates "Already viewed/job listed twice" as a reason to not save a vacancy, when they have already viewed the other artificial vacancy in the pair.

[29] In addition we provide robustness checks excluding observations of cases where individuals indicated "Already viewed/job listed twice" as a reason not to save a vacancy after having saved the first vacancy in the pair.

[30] Figure 8 contains the Edinburgh pairs only. Job openings in Glasgow received very little attention, only five times an artificial vacancy in Glasgow was viewed (see Figure **??** in the Appendix).

Figure 8: Number of times a vacancy was viewed

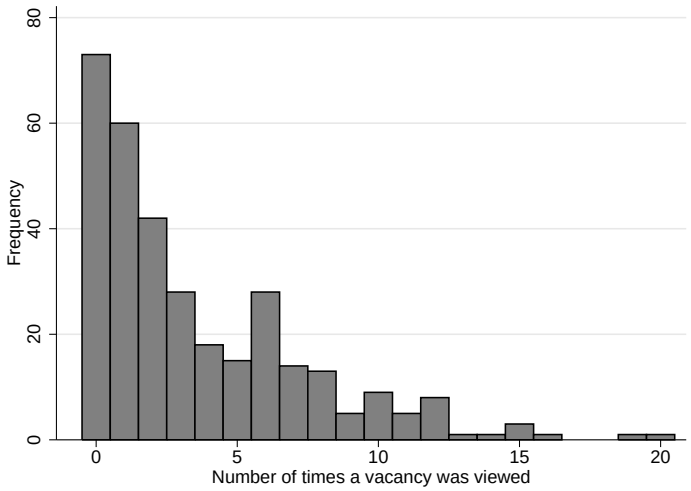

5.1 Do higher wages generate more interest in vacancies?

We start by exploring how the posted wage affects the number of interested job seekers. Since we are comparing the number of views between two vacancies within a pair, we perform a T-test for paired samples. The randomization process assigns one of the vacancies within a pair to either a wage increase or decrease. The other vacancy in the pair has the original wage.

Table 3: T-tests (paired) for mean number of views

Salary change	Baseline	Changed Wage	Difference	P-value	Obs. (Pairs)
Decrease	1.8	1.7	-0.1	.71	32
Increase	3.6	4.4	0.8	.00	131

We perform the test separately for pairs where the wage was increased and pairs where the wage was decreased. Results for the two t-tests are presented in Table 3. Column (1) shows the mean number of views for vacancies with the original wage, (2) shows mean number of views for changed wage vacancies, (3) shows the difference in means, (4) shows the p-value for a two-sided t-test for equal means and (5) shows the number of observations (pairs). We find that a decrease in the wage reduces the number of views, but the difference is not statistically significant. Note that the number of observations is small (32), because as mentioned earlier for several vacancies we could not decrease the wage due to the legal minimum wage level. For wage increases we have more observations (131) and we find that the number of views increases significantly, with 0.77 (21%). The p-value is smaller than 0.01. Note that this increase is related to a (weighted) average of

wage increases of 20% and 40%.[31]

To exploit the variation in the magnitude of the wage changes we perform a regression analysis. We regress the logarithm of the number of views on the percentage change in the wage (Δw) in order to estimate the elasticity.[32] To exploit the pair structure of the data, we include pair fixed effects (γ_p).

$$\ln V_{ip} = \alpha + \gamma_p + \beta \Delta w_{ip} + \epsilon_{ip} \tag{1}$$

Vacancies are indexed by subscript i. The parameter of interest is β, which measures the percentage effect of a one percent increase in the salary of a vacancy on the number of views. We use the logarithm of the number of views to take into account that the popularity of vacancies differs greatly (for example across occupational groups) and therefore a percentage effect is more appropriate that an constant effect. Note that most of the artificial vacancies were used twice: in the first and in the second wave. This implies that each pair typically has four artificial vacancies, where two have the original salary and the other two have the same altered wage.[33] Standard errors are clustered at the pair level.

Estimation results are presented in Table 4. In column (1) we find a highly significant positive effect of the wage change on the number of views. The estimate of 0.36 implies that for a 1% increase in the wage, the number of views increases with 0.43 %. Replacing the percentage change in the wage with dummies for a wage decrease or increase in column (2) shows that the wage effect is symmetric for increases and decreases, however due to lower the effect for decreases is not significant. Postal codes were varied within pairs of vacancies to make sure the pair was not identifiable. The postal codes were assigned independently of the wage variation. In columns (4) and (6) we control for postal code variation by including fixed effects for the outward code (the first three or four digits of the postal code). There are 14 of such areas in our dataset for which we have sufficient observations to include a fixed effect. As expected, we find that including these fixed effects does not influence the results in column (3) and (4). The outcome variable V_i is a count variable and has a mode of zero. An alternative approach is to use Poisson regression, which has been argued to be preferable over log-linear regression (e.g. Cameron and Trivedi (2005)). In column (5) we find that performing a Poisson regression leads to a slightly larger estimate. In column (6) we find that the stronger effect comes from wage increases, while the effect for decreases somewhat. Note that the number of observations is reduced somewhat as pairs with all zero views are dropped.[34]

For vacancies with a relatively low wage we could not reduce the wage and increased it instead. As result our sample contains relatively more wage increases among 'low wage' vacancies than among 'high wage' vacancies. This selection affects our estimates if one expects the effect of a wage change to be (i) different for increases and decreases and (ii) different for low and high wage vacancies. We present a number of estimates to address the potential impact of only assigning wage increases to low-wage vacancies in Table 5.

[31] Rather then using the number of views, we can also use the logarithm of the number of views as the outcome. The results are very similar and presented in Tables **??** and **??** in the Appendix.

[32] We take $\log V_i = log(\#\text{viewed}+1)$ to deal with the vacancies that were never viewed.

[33] Due to the restriction on the number of artificial vacancies, we posted somewhat less vacancies during the second wave. As a result not all artificial vacancies were used twice.

[34] In Table 15 the appendix we show that the results are almost identical when also Glasgow pairs are included in the analysis.

Table 4: Effect of wage change on number of views

	Log-linear regression				Poisson regression	
			postcode f.e.			
	(1)	(2)	(3)	(4)	(5)	(6)
Salary difference	0.36**		0.36***		0.54**	
	(0.16)		(0.13)		(0.21)	
Salary decrease		-0.10		-0.09		-0.07
		(0.09)		(0.09)		(0.17)
Salary increase		0.11**		0.11**		0.19***
		(0.05)		(0.05)		(0.07)
Pair fixed effects	yes	yes	yes	yes	yes	yes
Postal code f.e.	no	no	yes	yes	yes	yes
N	322	322	322	322	304	304

* p < 0.10, ** p < 0.05, *** p < 0.01. Robust standard errors in parentheses. The independent variable is log(viewed+1)in columns (1)-(4).

First, we show results using only wage increases in column (1). The resulting estimate on the salary change is almost identical to the estimate using both increases and decreases: 0.34 (compared to 0.36 in Table 4). Also the Poisson model estimate is almost the same (0.51, column (2)). Second, we show results using only wage increases and reweighing all low-wage vacancies (vacancies less than 40% above the minimum wage) by the share of high-wage vacancies that have a wage increase. The share of the high wage vacancies for which the wage was increased is 65%, such that we give low-wage vacancies a weight of 0.65. The result is presented in column (3) and is very similar to the unweighted estimate in column (1). Finally, we estimate a model with separate effects for (i) an increase for a low-wage vacancy (ii) an increase for a high wage vacancy and (iii) a decrease for a high-wage vacancy. The results in column (4) show that the positive effect of wage increases is stronger for low-wage vacancies. Wage increases for high wage vacancies increases the number of views, but not significantly. Wage decreases reduce the number of views, but also not significantly.

Rather than using views as the outcome variable, we can also consider saves. These are arguably a better proxy for applying to a vacancy, but may be affected by some job seekers distrusting the similarity between vacancies within a pair. We present analyses using saves as the outcome variable, analogous to those in Table 4 and 5 for views, in the Appendix in Tables 16 and 17. The results are very similar, with significantly positive coefficients for the salary difference effect on the number of saves. Also the magnitude is similar (the baseline estimate for the elasticity is 0.29 compared 0.36 for views).

5.2 Is the reservation wage property satisfied?

The second part of the empirical analysis will focus on whether the reservation wage property, which is an inherent feature of random search models but is violated in directed search models, holds in this context. To do this, we focus on individuals' decisions

Table 5: Effect of wage change on number of views: restricted samples

	(1) Increases	(2) Increases Poisson	(3) Increases Reweighted	(4) All
Salary difference	0.33**	0.51**	0.33**	
	(0.14)	(0.22)	(0.14)	
Salary increase*low				0.15**
				(0.07)
Salary increase*high				0.08
				(0.06)
Salary decrease*high				-0.09
				(0.09)
Pair fixed effects	yes	yes	yes	yes
Postal code fixed effects	yes	yes	yes	yes
N	258	248	258	322

* p < 0.10, ** p < 0.05, *** p < 0.01. Column (1), (3) and (4) are log-linear models, column (2) is a Poisson regression model. Column (1), (2) and (3) exclude vacancy pairs with a wage decrease. Column (4) contains increases and decreases.

regarding vacancies from the same pair. Again, we proxy applications by looking at viewings or savings. We include only Edinburgh pairs in the analysis (however including Glasgow pairs would make no difference as they were hardly ever viewed).

In particular we document a number of statistics relating to the relationship between viewing one or both of the vacancies in a pair. First, we will show the probability of not viewing the high-wage vacancy, conditional on viewing the low wage vacancy and viceversa:

$$P(V_h = 0|V_l = 1) \tag{2}$$
$$P(V_l = 0|V_h = 1) \tag{3}$$

Where $V_l = 1$ if the low wage vacancy in a pair was viewed, and similar for V_h. These probabilities are shown in column (1) of Table 6, in the first and second row. The number of observations are given in brackets. We find that out of all individuals that view the low wage vacancy in a pair, 38% does not view the high wage vacancy. The reverse probability, not viewing the low wage vacancy when the high wage vacancy is viewed, is 47%. If the reservation property holds, one would expect none of the individuals that view the low-wage vacancy to not also view the high wage vacancy.

One may worry that these probabilities do not fully represent conscious decisions of job seekers. For example, due to a large number of search results one of the vacancies in the pair may not appear on the first screen of results. If the job seeker does not continue browsing to the next page he may simply not observe the second vacancy.[35] The

[35]If, on the other hand, it occurs because one of the two vacancies does not fulfill the search criteria

Table 6: Viewing and saving probabilities

		conditional on listing both		
			high wage appearing first	low wage appearing first
	(1)	(2)	(3)	(4)
Viewing				
$P(V_h = 0 \mid V_l = 1)$	0.38 (520)	0.36 (500)	0.21 (200)	0.50 (300)
$P(V_l = 0 \mid V_h = 1)$	0.47 (623)	0.45 (590)	0.55 (344)	0.34 (246)
P-value test for equal proportions[a]	0.00	0.00	0.00[b]	0.20[b]
Saving				
$P(S_h = 0 \mid S_l = 1)$	0.42 (278)	0.39 (267)	0.16 (97)	0.54 (170)
$P(S_l = 0 \mid S_h = 1)$	0.49 (337)	0.48 (318)	0.61 (185)	0.36 (133)
P-value test for equal proportions[a]	0.08	0.03	0.00[b]	0.18[b]

Number of observations in brackets. All of the reported fractions in this table are significantly different from zero with p-value<0.001. [a] P-values from testing $P(V_h = 0|V_l = 1) = P(V_l = 0|V_h = 1)$ and similar for savings (both are two-sided tests). [b] Note that these are tests for cross probabilities (for column (3) H_0: 0.20=0.35 and for column (4) H_0: 0.50=0.55) and similar for savings (for column (3) H_0: 0.15=0.37 and for column (4) H_0: 0.54=0.61)

advantage of our experimental setup is that we can observe which vacancies a job seeker has 'listed' on his/her screen. So we can compute the above probabilities conditionally on listing both vacancies:

$$P(V_h = 0|V_l = 1, L_l = 1, L_h = 1) \quad (4)$$
$$P(V_l = 0|V_h = 1, L_l = 1, L_h = 1) \quad (5)$$

These probabilities are listed in the column (2) of Table 6. We find that the conditional probabilities of not viewing one of the two are only slightly lower for individuals that have listed both. Still in 36% of the cases that an individual views the low wage vacancy, (s)he does not view the high wage vacancy. Note however from the almost identical number of observations, that almost all participants that view one of the two vacancies have listed both vacancies.

We would like to exclude as much as possible that the decision to only view one of the two vacancies is caused by other factors. One such factor is the timing at which the vacancies were posted. In order to make the artificial vacancies less suspicious, the posting dates within the pair were varied slightly. The difference was never more than 2 days. Since by default search results were ordered by posting date, this prevented the two vacancies from appearing directly below each other. One may worry that job seekers

imposed by the job seeker, it is less of worry. In that case the choice as to view one of the two can be regarded as simply revealing preferences.

start at the top of the list when scanning vacancies, and therefore the ones that appear lower on the list are not viewed simply because the individual does not reach that part of the list. This reasoning is supported by the strong correlation that we observe between the posting date and the likelihood of viewing the vacancy. The earlier posted vacancies within a pair were, on average, viewed 2.46 times, while the later posted vacancies were viewed 3.53 times.

To address this concern, we compute probabilities similar to (4) and (5), but condition on which vacancy was posted first. Remember that posting dates were assigned at random, and are therefore independent of the wage. The results are presented in columns (3) and (4) of Table 6. We find that, as expected, the likelihood of viewing increases when the vacancy appears first (due to being posted more recently). For those that view the low wage vacancy while the high wage one appears first, 21% does not view the high one. This percentage is considerably lower than the 36% in column (2), however it is still a non-trivial share. The reverse probability (not viewing the low one when one views the high one and the low one appear first) is 34% (column (4)).

We perform two tests to assess the statistical significance of these probabilities. First, we test whether these probabilities of not viewing are significantly different from zero. A one-sample proportion test for each of the viewing probabilities shows that each of these is highly significantly different from zero (p-value always smaller than 0.001). Second we test for differences between viewing the high and low wage vacancies. In particular, we test whether $P(V_h = 0|V_l = 1) = P(V_l = 0|V_h = 1)$. Using a two-sample proportion test we find that the unconditional probabilities (column (1)) are significantly different (p-value<0.01). Also the probabilities conditional on listing both (column (2)) are significantly different (p-value<0.01). When conditioning on the posting order, we find that the probability of not saving the high one when the low one appears first (0.50) is not significantly different from the reverse probability (0.55) (p-value=0.20). However, the probability of not saving the high one when the high one appears first (0.20) is significantly different from the reverse probability (0.35) (p-value<0.01). Overall, these findings are in line with the results in subsection 5.1: a higher wage significantly increases the interest expressed by job seekers.

As mentioned in subsection 4.3, we can compute similar probabilities for saving vacancies, rather than only viewing them. These probabilities are reported on the third and fourth row of Table 6. The probability of not saving a high wage vacancy when one has saved the low wage vacancy is 0.42, while it is 0.39 when conditioning on listing both. If, in addition, the high wage vacancy appears first on the list, this probability becomes 0.16. All of these are quite close to the corresponding viewing probabilities.

5.3 Robustness checks

We perform a number of robustness checks to address certain concerns in our analysis. First, we exclude individuals that save an artificial vacancy, then view the second vacancy in the pair and do not save it indicating "it is posted twice/already viewed" as the reason for not saving. Second, we exclude individuals that responded in an exit survey that they sometimes, often or always could identify artificial vacancies. Third, we exclude vacancy pairs for which the geographical distance within the pair was more than 1 kilometer. Fourth, we discuss possible learning effects among participants.

5.3.1 Posted twice/already viewed

One may worry about individuals that save an artificial vacancy, then view the second vacancy in the pair and do not save it indicating "it is posted twice/already viewed" as the reason for not saving. These individuals may have either identified the artificial pair or simply believe they already saw the second vacancy since it closely resembles the first one. In either case it provides an explanation for not saving the second one, which is different from the hypothesis we are testing. Note however that this can affect the saving decision, but not the viewing decision. To investigate whether these events drive our results, we remove all the viewings and savings of this type from our data. This implies removing 37 pair-viewings, or 74 vacancy-viewings (out of a total of 1145 vacancy-viewings). The new probabilities are presented in Table 7, which has the same structure as Table 6 in terms of conditioning. We find that removing the 74 observations has very little effect on the probabilities: the non-viewing probabilities increases slightly (row 1), while the non-saving probabilities decrease slightly (row 3).[36] We conclude that the small number of cases in which participants suspect the pair to be the same vacancy does not drive the result that a non-trivial share of job seekers is interested in the low wage vacancy but not in the high wage vacancy.

5.3.2 Identifying artificial vacancies

At the end of the study we performed a short exit survey, asking participants whether they felt they could identify the artificial vacancies. The responses were: 'Never' (48%), 'Rarely' (20%) 'Sometimes' (28%) 'Often' (6%) 'Always' (1%). Participants that (believe they) can identify the artificial vacancies might change their behavior accordingly. To assess whether this is the case we redo our analysis including only those that responded 'Never' or 'Rarely'.[37] This implies dropping 654 vacancies viewings, or 57% of the ob-

[36] This is as expected: we removed observations in which both vacancies were viewed, such that the probability of only viewing one of the two must increase. On the other hand, in the removed observations only one of the two vacancies was saved, such that the probabilities of saving only one of the two must decrease.

[37] Note that since the exit survey was performed after the last session, we only have responses for those who participated in the final session (about 50%). Therefore we also drop all participants that did not

Table 7: Viewing and saving probabilities excluding participants that indicate "it is posted twice/already viewed" after having saved the first vacancy

		conditional on listing both		
			high wage appearing first	low wage appearing first
	(1)	(2)	(3)	(4)
$P(V_h = 0 \mid V_l = 1)$	0.40 (483)	0.38 (463)	0.22 (183)	0.53 (280)
$P(V_l = 0 \mid V_h = 1)$	0.50 (586)	0.48 (553)	0.58 (327)	0.37 (226)
$P(S_h = 0 \mid S_l = 1)$	0.38 (263)	0.35 (252)	0.15 (96)	0.49 (156)
$P(S_l = 0 \mid S_h = 1)$	0.46 (315)	0.45 (296)	0.56 (169)	0.34 (127)

Number of observations in brackets

Table 8: Viewing and saving probabilities including only participants that respond in exit survey that they could never or rarely identify artificial vacancies

		conditional on listing both		
			high wage appearing first	low wage appearing first
	(1)	(2)	(3)	(4)
$P(V_h = 0 \mid V_l = 1)$	0.38 (226)	0.36 (218)	0.16 (73)	0.52 (145)
$P(V_l = 0 \mid V_h = 1)$	0.49 (263)	0.45 (246)	0.53 (129)	0.36 (117)
$P(S_h = 0 \mid S_l = 1)$	0.36 (120)	0.34 (116)	0.06 (34)	0.52 (82)
$P(S_l = 0 \mid S_h = 1)$	0.51 (148)	0.48 (139)	0.57 (76)	0.37 (63)

Number of observations in brackets

servations. The resulting viewing and saving probabilities are presented in Table 8. The probabilities of not viewing or saving the high wage one (rows 1 and 3) are slightly lower. Only the probability of not saving the high one when it appeared higher is reduced quite a bit, to 0.06.

5.3.3 Geographical distance within a pair

Table 9: Viewing and saving probabilities including only pairs with a geographical distance below 1 kilometer

		conditional on listing both		
			high wage appearing first	low wage appearing first
	(1)	(2)	(3)	(4)
$P(V_h = 0 \mid V_l = 1)$	0.35 (264)	0.33 (256)	0.23 (96)	0.43 (160)
$P(V_l = 0 \mid V_h = 1)$	0.46 (327)	0.44 (306)	0.55 (176)	0.32 (130)
$P(S_h = 0 \mid S_l = 1)$	0.46 (139)	0.43 (135)	0.14 (41)	0.61 (94)
$P(S_l = 0 \mid S_h = 1)$	0.52 (171)	0.51 (158)	0.65 (90)	0.35 (68)

Number of observations in brackets

We attribute differences in interest for the vacancies within a pair to the wage differences. However, within a pair, the location was varied slightly in order to make sure the two vacancies would not appear to be the same. The postcode was changed for one of the two. In some cases this implied only a small difference in geographical location, however in other cases the difference can be larger. Both are however always posted in Edinburgh. Based on the 'outward code' (the first three or four digits of the post code), we compute the geographical distance between the two vacancies in the pair. In case this distance is small, it is unlikely to cause a difference in job seekers' interest. We find that in 45% of

finish the final survey.

the Edinburgh pairs the two vacancies share the same outward code. For 5% of the pairs the geographical distance is less than 1 kilometer, while for the other 50% the distance varies from 2 to 12 kilometers.[38] As a robustness check we perform the empirical analysis including only the 50% vacancy pairs for which the distance is less than 1 kilometer. Results are presented in Table 9. Again we find that our main results persist among pairs with little geographical distance. The probabilities are almost identical to those using the entire sample of vacancies.

5.3.4 Learning

A potential concern in our experimental design is learning among participants. Perhaps participants become more aware of artificial vacancies over the 12 weeks, or they become better at identifying the artificial vacancies. Such a learning process could be especially strong among those that encounter (several) artificial vacancies in the first weeks. To investigate whether learning occurs, we split the observations into three 4-week periods and check whether our empirical results differ between these periods. Both the number of participants and the number of fake vacancies decreased slightly over the 12 weeks, such that absolute numbers are uninformative. Alternatively we compute the ratio of savings per view. The mean of this ratio is 0.49 in period 1 (weeks 1-4), 0.41 in period 2 (weeks 5-8) and 0.59 in period 3 (weeks 9-12). Two-sided t-tests for equality show that the difference between the first and second period is not significant (p-value 0.12), while the difference between period 2 and 3 is significant (p-value<0.01). The difference between period 1 and 3 is borderline significant (p-value 0.09). Even though these results suggest differences between the periods, there is no monotone trend and the saving rate is actually highest in the last period. Such a pattern is difficult to reconcile with any plausible learning story. Rather we attribute it to differences in the pool of artificial vacancies and differences in the pool of participants (due to attrition). Furthermore, we have shown in subsection 5.3.2 that only including individuals that indicated they were unable to identify artificial vacancies makes little difference for the results.

5.4 Perceptions about competition and working conditions

Our results suggest that some job seekers prefer to apply to a low-wage job over a high-wage, when all other characteristics are equal. This result can be due to higher expected competition for the job at the high wage vacancy, as directed search models predict. However, the higher wage might also be interpreted as a signal of worse working conditions. If job seekers assume that higher wages are only offered to compensate for differences in working conditions, it is rational (for some) to prefer the low wage job.

To investigate whether this is the case, we performed an online survey. Each survey participant would be presented with a random set of 20 artificial vacancies, with 3 questions about each vacancy. The questions were as follows:[39]

1. How good would you expect an applicant needs to be in order to be considered for this job?

[38] Note that these are straight-line differences and the corresponding travel distances will be somewhat larger.

[39] Each question was multiple choice with options: Very much above average; Above average; Average; Below average; Very much below average.

2. How much competition would you expect for this job relative to other jobs in the same profession and area?

3. How would you expect the overall (non-monetary) working conditions of this job to be?

Each artificial vacancy was presented to at least 5 participants. The 20 vacancies that each participant was presented with, always consisted of 10 pairs of vacancies, such that we can observe within-individual variation in the answers about the two vacancies that only differ in salary. The different sets of 20 vacancies very fixed, but the order in which the vacancies were shown within a set was randomized for each participant. A total of 121 participants took part in the survey, and we have between 5 and 9 responses for each vacancy.[40]

Participants were recruited through the behavioural lab of the University of Edinburgh. They received a £20 internet voucher in return for completing the survey. To incentivize truthful answers as well as sufficient effort, we offered an additional £20 voucher to the participant in each vacancy set whose answers were closest to the average response. In order to receive reliable perceptions about vacancies, it is preferable to have participants that have some experience in the labor market. Therefore we tried to minimize the number of participating students by excluding university email addresses from the invitation email list. The resulting sample of participants has an average age of 31 (with a range from 19 to 62). 74% is employed (parttime of fulltime), 14% is unemployed and looking for work and 6% is self employed. Of all participants, 56% indicates to be currently looking for work (even if employed). The participants contained more women (68%) than men (32%).

Since the responses are ordinal, and each participant assessed both vacancies in a pair, we test for differences between high and low wage vacancies using the Wilcoxon signed rank test. Results are presented in Table 10. For all three questions, the responses are significantly higher for the high-wage vacancies. This implies that vacancies with a higher wage (but the same in all other respects) are perceived to (1) require applicants to be of higher quality, (2) attract more competition for the job and (3) have better non-monetary working conditions. These findings are in line with a directed search interpretation of our empirical findings: the high wage job is more attractive (both in monetary and non-monetary conditions), but is expected to attract more applicants and will thus be harder to get. Note that this rejects an alternative hypothesis that high wages compensate for worse working conditions (compensating differentials), which would have led to a reverse result on question 3. Thus, the survey results suggests that lack of interest in high-wage vacancies should be caused by higher expected competition for the job.[41]

We investigate whether the differences in perceptions about vacancies drive our result that high wage vacancies attract more interest. To do so, we compute the average responses across participants for each vacancy, using linear values ranging 1-5 for the re-

[40] Only the Edinburgh-pair vacancies were included in the survey.

[41] To assess the quality of the measurements of perceptions about the vacancies, we compute Cronbach's alpha for each set of individuals that judge the same 20 vacancies. There are 16 of such sets. The distribution of the 16 values is presented in the appendix in panel (a) of Figure 13. We find that 12 of the alpha's are above 0.7 and 15 are above 0.6, suggesting high agreement of the different measures. Computing the alpha's separately for the three questions (panels (b)-(d) in Figure 13), we find that agreement is particularly high for question 1 and 3, while somewhat lower for question 2.

Table 10: Wilcoxon paired signed-rank test comparing

	High wage response higher	Low wage response higher	Equal Response	p-value	Observations
Question 1	379	211	620	<0.001	1210
Question 2	479	230	501	<0.001	1210
Question 3	374	221	615	<0.001	1210

sponse choices. Next, we include the mean responses as control variables when regressing log-viewed on the salary difference (with a specification equal to that used in section 5.1).

The results are presented in columns (2)-(5) in Table 11. The perceptions are never significant and including them leads to only minor changes in the salary difference coefficient. To control for heterogeneity across survey participants we compute for each survey participants the average response to each of the three questions and subtract this average from his/her responses. The resulting deviations from the individual mean are then used to computed average responses for each vacancy and included in the regressions. The results are presented in columns (6)-(9) in Table 11. These demeaned variables are also never significantly different from zero and lead to only minor changes in the salary difference coefficient. We perform the same exercise with the number of saves instead of views and present results in the appendix in Tables 18. Also in this case the perceptions are never significant and only change the salary difference coefficient slightly.

6 Directed search model

In the empirical section we documented two facts. First, a job that offers a high wage attracts more applicants than the same job with a lower wage. Second, a non-trivial fraction of applicants applies to the low wage job, but not to the high-wage job even if they observe both jobs (i.e., both are listed on their screen). We discussed that these findings are qualitatively in line with predictions from directed search. It remains an open question whether the magnitudes that we find are plausible as well. To this end, we calibrate a directed search model to UK data.

Our attempt is to use the simplest model that still allows us to meaningfully talk about both of the above facts. Despite our aim for parsimony, we need to include a minimum number of modelling elements in order to achieve this. We discuss these in turn to highlight exactly the role of each of them, and proceed to the formalities thereafter. In order to assess the second fact we need a model with multiple job applications. Otherwise it follows trivially that a worker who applies to the low wage will not apply to the high wage, since he has no further application left, and it is impossible to match the fact that some workers do apply to both wages. Among the existing models with multiple job applications by Albrecht et al. (2006), Galenianos and Kircher (2009), Kircher (2009) and Wolthoff (2014) it turns out to be convenient to work with one that retains a single offered wage in equilibrium, as in one-application models with homogeneous agents. The reason is technical: for tractability nearly all directed search models use a continuum of vacancies that are observed by a continuum of workers. Fixing a given worker who applies to a low wage job, and fixing a high wage job, the chances that this worker applies

Table 11: Effect of salary on views, controlling for perceptions

	(1) logviewed	(2) logviewed	(3) logviewed	(4) logviewed	(5) logviewed	(6) logviewed	(7) logviewed	(8) logviewed	(9) logviewed
Salary difference	0.36***	0.32**	0.46**	0.26*	0.36**	0.29*	0.43**	0.25	0.34*
	(0.13)	(0.15)	(0.18)	(0.15)	(0.18)	(0.15)	(0.17)	(0.15)	(0.18)
Question 1		-0.05			-0.04				
		(0.12)			(0.13)				
Question 2			0.09		0.15				
			(0.10)		(0.11)				
Question 3				-0.13	-0.17				
				(0.11)	(0.13)				
Q1 dem.						-0.09			-0.06
						(0.12)			(0.13)
Q2 dem.							0.06		0.14
							(0.10)		(0.11)
Q3 dem.								-0.14	-0.18
								(0.11)	(0.13)
Constant	1.07***	1.21***	0.83**	1.48***	1.34***	1.07***	1.06***	1.07***	1.06***
	(0.18)	(0.35)	(0.34)	(0.40)	(0.47)	(0.18)	(0.18)	(0.19)	(0.18)
N	322	322	322	322	322	322	322	322	322

Standard errors in parentheses

* p¡0.10, ** p¡0.05, *** p¡0.01

there is zero (since any job attracts only a zero measure of workers). This is different if one assumes that workers observe only a finite number of jobs. It turns out that finite observability does not alter at all the equilibrium of a directed search model with a unique equilibrium wage (see Acemoglu and Shimer (1999) for a first version of this insight). We will therefore use such a model and consider the reaction to a deviation away from the equilibrium wage, which allows us to talk about the points above. Only the model by Albrecht et al. (2006) - in which firms bid up the wage in case multiple of them pursue the same worker - has such a structure.

Finally, it is hopeless to address the first fact in a model that has no on-the-job search, as in its absence the starting wage determines the wage throughout the employment spell and workers react excessively to higher wage offers. This is the same reason why random search models without on-the-job search fail to match wage dispersion for comparable workers by orders of magnitude (Hornstein, Krusell and Violante, 2011), as workers again react too much to the presence of high wage jobs and would not want to accept low wage jobs. With on-the-job search the starting wage is an imperfect predictor of the wage during the employment spell, as subsequent search can improve the wage, making workers less sensitive to the initial wage. We allow for on-the-job search, and for consistency again allow firms to bid for the workers' services in case another firm is also interested in the same work, which places the structure of Postel-Vinay and Robin (2005?) into a directed search setting. Combining multiple simultaneous job applications with subsequent on-the-job search is new, but uses the existing elements from the literature and is necessary to assess the facts.

6.1 The model

We lay out a simple directed job search model with multiple-applications based on Albrecht et al. (2006) (AGV). On-the-job search in introduced in an easy way: firms observe whether an applicant already has a job or not, and always offer the job first to someone who is unemployed. The logic is the following: employed workers use another job offer to bid up their wage, and the firm makes zero profit from them. Such workers would constitute an "adverse selection" problem if the firm would not observe them. Since the firm observes them, they only consider such workers if they do not have any other worker with whom they could possibly make a profit. As a result, unemployed workers are not affected by the on-the-job search of other workers since they would always get the offer first.

Assume time is discrete, β is the discount factor and workers have N applications per time period. Let $\lambda(w)$ be the number of applications from unemployed workers per job offering wage w. Assume that a random shock determines whether such an application is suitable or not. With probability $1-A$ an application is not suitable. This is immediately visible to the firm, and only the fraction A of applications is considered by firms. From the unemployed worker's perspective the probability of getting a job offer when applying to wage w is

$$m(\lambda(w)) = \frac{A(1-e^{-A\lambda(w)})}{A\lambda(w)} \tag{6}$$

$A\lambda(w)$ is the queue of other eligible applications, $(1-e^{-A\lambda(w)})/A\lambda(w)$ is the chance that the worker gets the offer if his application is eligible, and the multiplication of the

numerator by A takes into account that only with this probability his application is considered.

If the worker only get's this one job offer, he is paid the announced wage for the duration of the match (unless the wage is too low so that the worker prefers to remain unemployed). If he get's more than one offer he is paid his marginal product y for the duration of the match as firms bid up the wage to attract the worker. If a worker does not get any job this period, he get's unemployment benefits b. Once employed, there is a chance $1-\delta$ of loosing the job between one period and the next. We assume that a worker who applies today only starts on the job next period (and the separation shock only starts once he is actually on the job).

There is free entry: firm's pay entry cost c for posting a vacancy this period. They post a vacancy at their desired wage w. AGV show that in equilibrium there is only one wage w^*. If all other firms offer this wage, an individual firm maximizes

$$\max_{w\in[r,y]} n(\lambda(w))[1-m(\lambda(w^*))]^{N-1}\beta\Pi(w)-c$$

where $\Pi(w)$ is the net present value of having the worker until he separates and $n(\lambda(w)) = (1-e^{-A\lambda(w)})$.[42] Note that production starts next period, therefore the discounting. The firm only get's value from the worker if the other N-1 applications from this worker are not successful, otherwise it bids away all surplus. The firm cannot offer a wage below the reservation wage r of the worker, as otherwise the worker would not accept the wage even if this were the only offer. Let W be the value of this program. In equilibrium $W=0$.

The net present value (NPV) of having a worker is:

$$\begin{aligned}\Pi(w) &= y-w+\beta\delta\Pi(w)\\ &\Leftrightarrow \quad \Pi(w)=\frac{y-w}{1-\beta\delta}.\end{aligned}$$

AGV show in a model without on-the-job search that the equilibrium offered wage w^* falls to the worker's reservation value r at which the worker is exactly indifferent between accepting the job and not accepting (they apply nevertheless because they hope to get two offers, in which case they can bid up the wage). So in equilibrium the market utility of workers is determined by sending all N applications to firms offering $w^*=r$. Using $m^*=m(\lambda(w^*))$ to denote the chance that in equilibrium an application by the worker yields a job offer, the worker's equilibrium utility is:

$$\begin{aligned}U^* &= b+Nm^*[1-m^*]^{N-1}\beta V(w^*)\\ &\quad +[1-m^*]^N\beta U^*\\ &\quad +[1-Nm^*[1-m^*]^{N-1}-(1-m^*)^N]\beta V(y)\end{aligned}$$

It comprises in the first line the current payoff b and the NPV of working at the equilibrium wage w^* from next period onwards, which only happens if one of the N applications is successful but none of the others. The second line captures the possibility that none of the applications is successful, in which case the worker remains unemployed. In all other

[42] Note that applications from employed workers do not show up here as they are only getting an offer if no eligible unemployed worker is present, but then they generate no profit as they bid up the wage.

cases, captured by the last line, the worker gets to work at his marginal product y. Since $w^* = r$ is exactly the value where $V(r) = U^*$, and using $m^* = m(\lambda(w^*))$, this reduces to

$$\begin{aligned} U^* &= b + [1-m^*]^{N-1}[1+(N-1)m^*]\beta U^* \\ &\quad +[1-[1-m^*]^{N-1}[1+(N-1)m^*]]\beta V(y) \\ \Leftrightarrow U^* &= \frac{b + \left[1-[1-m^*]^{N-1}[1+(N-1)m^*]\right]\beta V(y)}{1-[1-m^*]^{N-1}[1+(N-1)m^*]\beta} \end{aligned}$$

where $V(.)$ is defined recursively as the value from having the job. If the worker works already at productivity y, his value is given as if he continued always at this firm (there is no need to search further, but if he does there are no further gains). The value function is given by

$$\begin{aligned} V(y) &= y + \beta\delta V(y) + \beta(1-\delta)U^* \\ &\Leftrightarrow \quad V(y) = \frac{y + \beta(1-\delta)U^*}{1-\delta\beta}. \end{aligned}$$

If the worker currently works at a job with wage $w = w^*$ then he continues searching. Assume an employed worker has $\tilde{N}$ applications and a matching efficiency $\tilde{A}$. It is easily possible that $\tilde{N} < N$ since employed workers have less time, but that $\tilde{A} > A$ if employment provides skills that make it easier for a worker to do other jobs.[43]

A worker employed at wage $w < y$ gets the wage this period, but if he does not loose the job he has the chance to move to wage y next period if he gets another job offer (with probability $\tilde{m}^*$). If not he continues earning wage w. If he looses the job he moves to unemployment. His value function is:

$$\begin{aligned} V(w) &= w + \beta\delta\tilde{m}^* V(y) + \beta\delta(1-\tilde{m}^*)V(w) + \beta(1-\delta)U^* \\ &\Leftrightarrow \quad V(w) = \frac{w + \beta\delta\tilde{m}^* V(y) + \beta(1-\delta)U^*}{1-\delta\beta(1-\tilde{m}^*)}. \end{aligned}$$

And so the value of unemployment is

$$\begin{aligned} U^* &= \frac{b + \left[1-[1-m^*]^{N-1}[1+(N-1)m^*]\right]\beta\frac{y+\beta(1-\delta)U^*}{1-\delta\beta}}{1-[1-m^*]^{N-1}[1+(N-1)m^*]\beta} \\ \Leftrightarrow U^* &= \frac{b(1-\delta\beta) + (1-Z)\beta y}{\beta^2(1-\delta-Z) + \beta(\delta+Z) - 1} \end{aligned}$$

With $Z = [1-m^*]^{N-1}[1+(N-1)m^*]$. This expression gives the value from unemployment as a function of the application success rate m^*, the discount factor β, the number of

[43] For workers employed at a wage below productivity the matching rate is given by $\tilde{m}^* = 1 - \left(1 - e^{-A\lambda(w)}\tilde{A}\frac{1-e^{-\tilde{A}\mu(w^*)}}{\tilde{A}\mu(w^*)}\right)^{\tilde{N}}$. This expression accounts for the fact that an application is only success if no unemployed worker applies (prob $e^{-A\lambda(w)}$). In this case the application is successful if it is eligible (with probability $\tilde{A}$) and gets an offer (with probability $\frac{1-e^{-\tilde{A}\mu(w^*)}}{\tilde{A}\mu(w^*)}$) where $\tilde{A}\mu(w^*)$ is the queue length of eligible applications by employed job seekers to vacancies). With complementary probability an application fails, and raised to power $\tilde{N}$ means that all of them fail. One minus this gives the probability that at least one application is successful.

applications N, the job loss rate δ, and the productivity which we can normalize to $y = 1$. Given the expression for U^*, the equilibrium wage $w^* = r$ can be recovered via:

$$\begin{aligned} U^* &= V(r) \\ U^* &= \frac{r + \beta\delta\tilde{m}^* V(y) + \beta(1-\delta)U^*}{1 - \delta\beta(1-\tilde{m}^*)} \\ (1 - \delta\beta(1-\tilde{m}^*))U^* &= r + \beta\delta\tilde{m}^* \frac{y + \beta(1-\delta)U^*}{1-\delta\beta} + \beta(1-\delta)U^* \\ r &= \frac{[(1-\delta\beta)(1-\beta) + \delta\beta(1-\beta)\tilde{m}^*]U^* - \beta\delta\tilde{m}^* y}{(1-\delta\beta)} \end{aligned}$$

Without on the job search ($\tilde{m}^* = 0$) this reduces to the familiar $(1-\beta)U^* = r$, so the reservation wage is equal to the flow value of unemployment. The reason why it is not that simple when there is on the job search is that benefits of a job not only includes the wage but also the option value of obtaining better jobs in the future. Note that the wage could be negative, in which case further thoughts are warranted.

6.2 Deviations

Now we proceed to deviations. Note that again only the behavior of unemployed workers is important. If a firm deviates by offering a wage $w > w^*$, the unemployed applicant's return is:

$$\begin{aligned} U(w) &= b + m(\lambda(w))[1-m^*]^{N-1}\beta V(w) \\ &\quad (N-1)m^*[1-m(\lambda(w))][1-m^*]^{N-2}\beta V(w^*) \\ &\quad +[1-m(\lambda(w))][1-m^*]^{N-1}\beta U^* \\ &\quad + \begin{bmatrix} 1-m(\lambda(w))[1-m^*]^{N-1} \\ -(N-1)m^*[1-m(\lambda(w))][1-m^*]^{N-2} \\ -[1-m(\lambda(w))][1-m^*]^{N-1} \end{bmatrix} \beta V(y) \end{aligned} \tag{7}$$

The logic is identical to before, only now we have to separately account for the deviation wage and the regular equilibrium wages. Those who apply to the new wage cannot make more than the market utility, so we have $U(w) = U^*$. Also, $V(w^*)$ equals U^* (if we do not impose a binding floor at a zero wage). Applying the implicit function theorem implies that the derivative of $\lambda(w)$ evaluated at w^* is

$$\lambda'(w^*) = \frac{[1-m^*]\frac{1}{1-\delta\beta(1-\tilde{m}^*)}}{-m'(\lambda(w^*))(N-1)\left[V(y) - U^*\right]}$$

And the elasticity is given by :

$$\frac{w^*\lambda'(w^*)}{\lambda(w^*)}$$

6.3 Calibration

We calibrate the model using statistics from the UK, where possible for the 4th quarter of 2013, which is the start of our experimental study. All values are listed in Table 12. We

Table 12: Calibration values

Parameter	meaning	value	source
Period length		Month	
β	Discount factor	0.9914	10% discounting per year
δ	1-Separation rate	0.9892	Sum of (UK, 1996-2010) job-to-unemployment hazard and job-to-inactivity hazard (which is 3.2% quarterly) (Gomes (2012))
y	Productivity	1	Normalized
N	Applications	12	Based on own study
Replacement rate	Ratio of benefits to average wage (in the population)	0.4	UK replacement ranges from 10 to 78% (depending on previous wage and family characteristics)
v	Vacancies	573,000	ONS statistics: vacancies in UK (2013,Q4)
u	Unemployment	2,348,000	ONS statistics: unemployed in UK (2013,Q4)
j	job finding rate	0.199	NOMIS statistics on off-flow of JSA claimants leads to 5.85% weekly (UK, Q4 2013), on the other hand Gomes (2012) reports a weekly unemployment-to-job hazard of 2.4%. We pick 5% weekly.
$\tilde{M}$	job-to-job transition rate	0.0082	ONS statistics: 722,000 job-to-job transitions quarterly (UK, Q4 2013) while employment was 30,288,000
ρ	probability of switching jobs conditional on an offer	0.5	

take the time period to be a month, and set the number of applications $N = 12$ based on an observed average of 3 applications per week in our study. We set the discount factor such that there is 10% discounting per year. Market tightness v/u follows from the UK Office for National Statistics (ONS), which, together with N, pins down the equilibrium queue length $\lambda(w^*)$.

According to NOMIS statistics, the off-flow of job seekers' allowance claimants implies a 5.85% weekly job finding rate in the UK in the fourth quarter of 2013. This might be an overestimate if non-claiming job seekers have a lower job finding rate. Alternatively, Gomes (2012) reports a weekly unemployment-to-job hazard of 2.4%, but this figure is an UK average over 1996-2010. We pick an intermediate value of 5% per week. Since the job finding rate equals $1 - (1 - m^*)^N$, this determines the equilibrium job offer probability m^*. Given m^* and $\lambda(w^*)$ we can back out the value of A from equation (6).

For the separation rate we use statistics from Gomes (2012) and take the sum of

Table 13: Model predictions

	Model calibration	Experimental result
Elasticity	0.49	0.36
$P(A_h = 0 \mid A_l = 1)$	0.47	0.41
$P(A_l = 0 \mid A_h = 1)$	0.60	0.51

the job-to-unemployment hazard and the job-to-inactivity hazard. The corresponding monthly separation rate is 0.0108, such that $\delta = 0.9892$. Productivity y is normalized to 1. For unemployment benefits we target a replacement rate of 40% of the average wage in the population and adjust b accordingly. The UK replacement rate is highly dependent on family and job characteristics and varies from 10 to 78% (based on OECD Tax-benefit models).

Job-to-job transitions are somewhat more complex. Only the share of the workers with wage $w = r$ searches on the job, because there are no gains to moving jobs for those that already work at $w = y$. Furthermore, a worker receiving a competing job offer has his wage bid up to productivity, but it is unclear whether he moves to the new firm or receives the wage increase at his current job. We define this probability as γ and set it equal to 0.5. The job-to-job transition rate observed in the data ($\tilde{M}$) then equals $\gamma n_w \tilde{m}/(n_w + n_y)$, where n_w is the share of workers earning the low wage and n_y the share of workers earning the high wage. The job-to-job transition rate in the UK in the fourth quarter of 2013 is 0.19% per week. Using steady state conditions for n_w and n_y we can back out the matching rate for employed workers $\tilde{m}$.

Now all elements necessary to compute r, w^* and $\lambda'(w^*)$ are known.[44] The elasticity $\lambda'(w^*)$ is presented in the first row of Table 13. Given our calibration, it is 0.49, which is reasonably close to the corresponding empirical value that we estimated in section 5, which is 0.36 (presented in the first column of Table 4).

The second outcome of interest is the probability of not applying to the high wage vacancy conditional on applying to the low wage vacancy: $P(A_h = 0 \mid A_l = 1)$. Calculating this probability requires setting an additional parameter X, equal to the number of vacancies that each job seekers observes. As discussed, it is required that each worker observes a random finite set of vacancies (rather than the continuum of vacancies), because otherwise the probability of applying to one particular vacancy would always be zero. In our study, we observe that individuals look at 25 vacancies per week, thus 100 per month. We conjecture that approximately one-third is somewhat interesting and set $X = 30$. Consider a vacancy that offers a wage above the equilibrium wage w^* (as was done in the manipulated vacancies used in the experiment). Denote the deviant wage as w and let p be the probability that an individual applies to this particular vacancy. The queue length at the deviant vacancy is given by p multiplied by the number of people that observe it:

$$\begin{aligned} \lambda(w) &= p\frac{Xu}{v} \\ p &= \frac{\lambda(w)v}{Xu} \end{aligned}$$

[44] The value of b is dependent on the shares of low and high wage earners in the model, which in turn are affected by b. By iterating the model until convergence we find $b = 0.12$.

Using that $\lambda(w^*) = N\frac{u}{v}$ we get:

$$p = \frac{\lambda(w)}{\lambda(w^*)}\frac{N}{X}$$

The value p should be such that wokers are indifferent between the deviant vacancy and equilibrium wage vacancies. Using equation 7 we compute $\lambda(w)$ and are thus able to calculate p. As a result, each worker has in expectation $N - p$ applications left for equilibrium wage vacancies, of which they observe $X-1$. Thus, the probability of applying at each of those is:

$$q = \frac{N-p}{X-1} \tag{8}$$

The probability of applying low, but not high is $q(1-p)$ and the probability of applying high but not low is $(1-q)p$. Using a wage difference of 20 % for the deviant vacancy, we find that these are 0.48 and 0.60, respectively (see Table 13. The corresponding probabilities obtained from the experiment are 0.41 and 0.51 (using saved vacancies as a proxy for applications, see Table 6). Thus, they are quite similar. We conclude that using a reasonable calibration, the model produces results that are of similar magnitude as those in our experiment.

7 Conclusion

In this study we present results on how the wage announcement in vacancies affects behavior of potential applicants. By posting pairs of vacancies with randomly assigned wages we provide evidence that, holding all vacancy characteristics constant, a higher wage attracts more applicants. Furthermore, we find support for one key prediction of the directed search model: some applicants only show interest in a low wage vacancy even though a high-wage vacancy, which is otherwise identical, exists. This finding is inconsistent with a model in which workers apply to each vacancy with an offer that exceeds the reservation wage. It suggests that applications are costly, and that workers take into account that high wage vacancies receive more applicants and thus constitute a lower probability of a job offer.

The strength of our approach lies in the realistic setting on which the observations are based. The participants were unemployed job seekers performing their usual job search, while the set of vacancies contained up 80% of all job openings in the UK. This ensures that the behavior observed in the study is likely to be very close to real life behavior. However, one caveat should be taken into account when interpreting the results. Our results are based on whether vacancies were saved, rather than whether actual applications were submitted. This limits effort and time spend by the participants on non-real vacancies, but it may not be a perfect predictor of actual applications.

We show the potential of applying the audit study methodology to the hiring side of the labor market. Such an approach can be used to investigate important questions regarding inequality in the labor. While extensive research has been done on firm behavior when selecting applicants, little is known about whether different types of workers apply to different vacancies. Such behavior could contribute to widely observed wage differentials. Unfortunately our study lacks the sample size to provide evidence on the composition

of applicants that opts for the lower-wage vacancies. We hope that by showing that our approach can be carried out with limited ethical concerns in terms of the burden on job seekers, we may contribute to further research on these topics.

References

Acemoglu, D. and Shimer, R. (1999). Holdups and efficiency with search frictions. *International Economic Review*, 40(4):827–849.

Albrecht, J., Gautier, P. A., and Vroman, S. (2006). Equilibrium directed search with multiple applications. *The Review of Economic Studies*, 73(4):869–891.

Belot, M., Kircher, P., and Muller, P. (2015). Providing advice to job seekers at low cost: An experimental study on on-line advice.

Bertrand, M. and Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review*, 94(4):991–1013.

Burdett, K. and Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 39(2):257–273.

Cameron, A. C. and Trivedi, P. K. (2005). *Microeconometrics: methods and applications*. Cambridge university press.

Channel 4 (2014). Why is government website carrying fake jobs? http://www.channel4.com/news/why-is-government-website-carrying-fake-jobs. Posted 07-02-2014. Last accessed 28-09-2015.

Computer Business Review (2014). Universal jobmatch here to stay despite fake job adverts. http://www.cbronline.com/news/cloud/aas/universal-jobmatch-here-to-stay-but-future-of-provider-monster-is-unclear-4204007. Written by Joe Curtis. Posted 26-03-2014; Last accessed 28-09-2015.

Dal Bó, E., Finan, F., and Rossi, M. A. (2013). Strengthening state capabilities: The role of financial incentives in the call to public service. *The Quarterly Journal of Economics*, 128(3):1169–1218.

Galenianos, M. and Kircher, P. (2009). Directed search with multiple job applications. *Journal of Economic Theory*, 144(2):445–471.

Gomes, P. (2012). Labour market flows: Facts from the united kingdom. *Labour Economics*, 19(2):165 – 175.

Holzer, H. J., Katz, L. F., and Krueger, A. B. (1991). Job queues and wages. *The Quarterly Journal of Economics*, 106(3):739–768.

Imai, K., King, G., and Nall, C. (2009). The essential role of pair matching in cluster-randomized experiments, with application to the mexican universal health insurance evaluation. *Statistical Science*, 24(1):29–53.

Joyce, S. P. (2015). How to avoid 5 major types of online job scams. http://www.job-hunt.org/onlinejobsearchguide/job-search-scams.shtml. Last accessed 28-09-2015.

Kircher, P. (2009). Efficiency of simultaneous search. *Journal of Political Economy*, 117(5):861–913.

Marinescu, I. (2014). The general equilibrium impacts of unemployment insurance: Evidence from a large online job board. Mimeo.

Marinescu, I. and Rathelot, R. (2014). Mismatch unemployment and the geography of job search. Mimeo.

Marinescu, I. and Wolthoff, R. (2014). Opening the black box of the matching function: The power of words. Mimeo.

Moen, E. R. (1997). Competitive search equilibrium. *Journal of Political Economy*, 105(2):pp. 385–411.

Neumark, D., Bank, R. J., and Van Nort, K. D. (1996). Sex discrimination in restaurant hiring: An audit study. *The Quarterly Journal of Economics*, 111(3):915–941.

The New York Times (2009a). Company rarely placed clients in jobs, former employees say. http://www.nytimes.com/2009/08/17/us/17careerbar.html. Written by Michael Luo. Posted 16-08-2009; Last accessed 28-09-2015.

The New York Times (2009b). Online scammers prey on the jobless. http://www.nytimes.com/2009/08/06/technology/personaltech/06basics.html. Written by Riva Richmond. Posted 05-08-2009; Last accessed 28-09-2015.

Wolthoff, R. (2014). It's about time: Implications of the period length in an equilibrium search model. *International Economic Review*, 55(3):839–867.

Appendix

Institutional Setting

In Figure 9 we present aggregate labor market statistics. Figure (a) shows the unemployment rate in the UK and Edinburgh since 2011. The vertical line indicates the start of our study. The unemployment rate in Edinburgh is slightly lower than the UK average, and is rather stable between 2011 and 2014. These statistics are based on the Labour Force Survey and not the entire population. Therefore we present the number of JSA claimants in the Edinburgh and the UK in panel (b), which is an administrative figure and should be strongly correlated with unemployment. We find that the number of JSA claimants is decreasing monotonically between 2012 and 2015, and that the Edinburgh and UK figures follow a very similar path. The number of JSA claimants in Edinburgh during our study is approximately 9,000, such that the 150 participants per wave in our study are about 2% of the stock. The monthly flow of new JSA claimants in Edinburgh during the study is around 1,800 (not shown in the graph).

Figure 9: Aggregate labor market statistics

(a) Unemployment rate

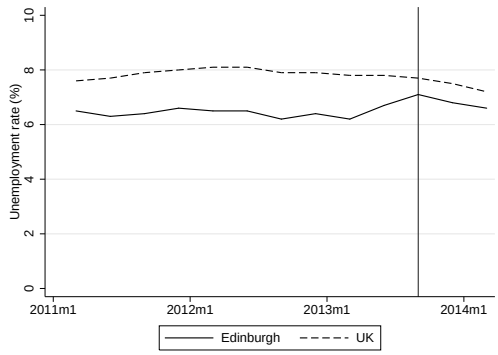

(b) JSA claimants

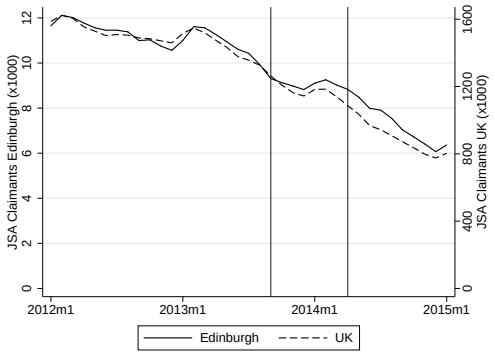

Additonal descriptive statistics

Figure 10: Number of times a vacancy was saved: Edinburgh pairs

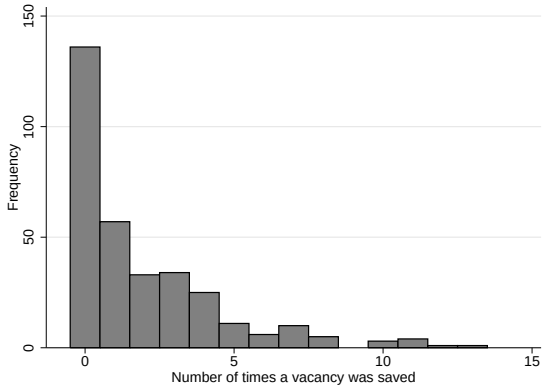

Figure 11: Number of times a vacancy was viewed/saved: Glasgow pairs

(a) Viewed

(b) Saved

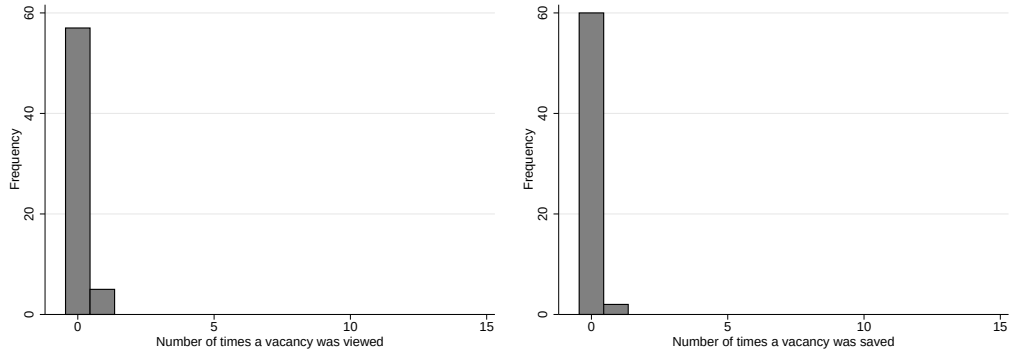

Figure 12: Salary differences of artificial vacancies

(a) Edinburgh pairs

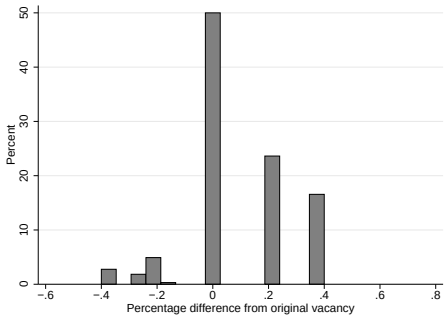

(b) Glasgow pairs

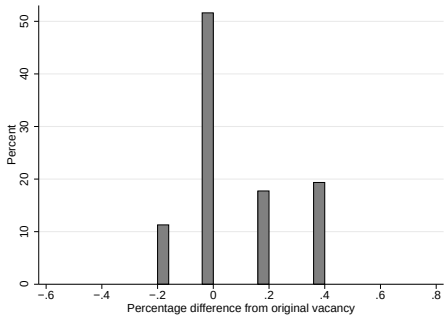

(c) Single vacancies

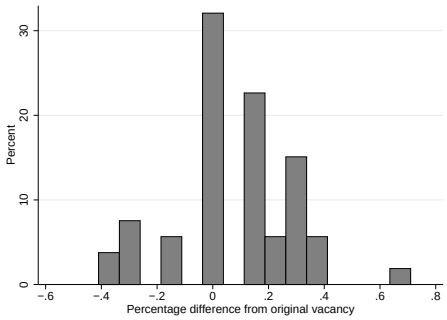

(d) Mixed pairs

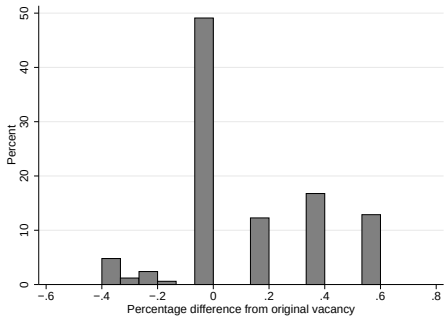

Additional empirical results

Table 14: T-tests (paired) for mean number of saves

Salary change	Baseline	Changed Wage	Difference	P-value	Obs. (Pairs)
Decrease	0.8	0.9	.1	.44	32
Increase	1.9	2.4	.5	.01	131

Table 15: Effect of wage change on number of viewed: Edinburgh pairs and Glasgow pairs

	Log-linear regression				Poisson regression	
			postcode f.e.			
	(1)	(2)	(3)	(4)	(5)	(6)
Salary difference	0.34***		0.34***		0.55***	
	(0.11)		(0.11)		(0.21)	
Salary decrease		-0.08		-0.07		-0.07
		(0.08)		(0.08)		(0.17)
Salary increase		0.11***		0.11***		0.20***
		(0.04)		(0.04)		(0.07)
Pair fixed effects	yes	yes	yes	yes	yes	yes
Postal code f.e.	no	no	yes	yes	yes	yes
N	384	384	384	384	318	318

* p < 0.10, ** p < 0.05, *** p < 0.01. Robust standard errors in parentheses. The independent variable is log(viewed+1) in columns (1)-(4).

Table 16: Effect of wage change on number of saved

	Log-linear regression				Poisson regression	
			postcode f.e.			
	(1)	(2)	(3)	(4)	(5)	(6)
Salary difference	0.29**		0.29**		0.52*	
	(0.12)		(0.12)		(0.27)	
Salary decrease		0.03		0.03		0.19
		(0.06)		(0.06)		(0.19)
Salary increase		0.13***		0.13***		0.22**
		(0.05)		(0.05)		(0.10)
Pair fixed effects	yes	yes	yes	yes	yes	yes
Postal code f.e.	no	no	yes	yes	yes	yes
N	322	322	322	322	240	240

* p < 0.10, ** p < 0.05, *** p < 0.01. Robust standard errors in parentheses. The independent variable is log(saved+1)in columns (1)-(4).

Table 17: Effect of wage change on number of saved: restricted samples

	(1) Increases	(2) Increases Poisson	(3) Increases Reweighted	(4) All
Salary difference	0.37**	0.54**	0.37**	
	(0.15)	(0.27)	(0.15)	
Salary increase*low				0.21***
				(0.08)
Salary increase*high				0.04
				(0.05)
Salary decrease*high				0.03
				(0.06)
Pair fixed effects	yes	yes	yes	yes
Postal code fixed effects	yes	yes	yes	yes
N	258	206	258	322

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses. Column (1), (3) and (4) are log-linear models, column (2) is a Poisson regression model. Column (1), (2) and (3) exclude vacancy pairs with a wage decrease. Column (4) contains increases and decreases.

Table 18: Effect of salary on saved, controlling for perceptions

	(1) logsaved	(2) logsaved	(3) logsaved	(4) logsaved	(5) logsaved	(6) logsaved	(7) logsaved	(8) logsaved	(9) logsaved
Salary difference	0.29**	0.35**	0.25	0.22	0.26	0.33**	0.25	0.21	0.26
	(0.12)	(0.16)	(0.16)	(0.14)	(0.18)	(0.16)	(0.17)	(0.14)	(0.18)
Question 1		0.07			0.12				
		(0.12)			(0.12)				
Question 2			-0.04		-0.03				
			(0.09)		(0.09)				
Question 3				-0.10	-0.12				
				(0.09)	(0.10)				
Q1 dem.						0.06			0.12
						(0.11)			(0.11)
Q2 dem.							-0.04		-0.02
							(0.09)		(0.10)
Q3 dem.								-0.10	-0.14
								(0.09)	(0.09)
Constant	0.41***	0.20	0.51	0.71**	0.52	0.41***	0.41***	0.41***	0.41***
	(0.13)	(0.36)	(0.31)	(0.33)	(0.47)	(0.13)	(0.13)	(0.14)	(0.14)
N	322	322	322	322	322	322	322	322	322

Standard errors in parentheses

* p¡0.10, ** p¡0.05, *** p¡0.01

Figure 13: Cronbach alpha's for the 16 vacancy sets

(a) All questions

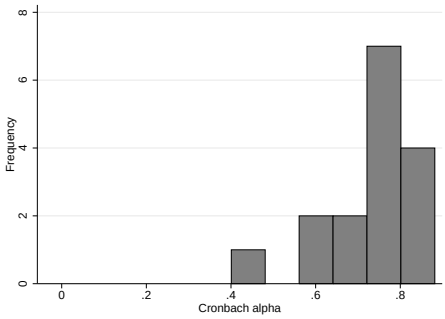

(b) Question 1

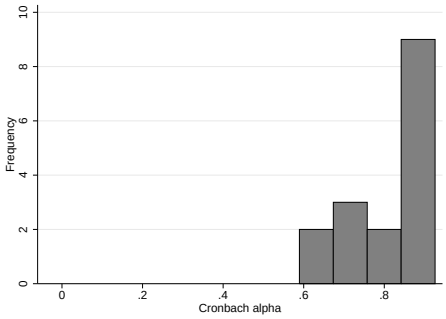

(c) Question 2

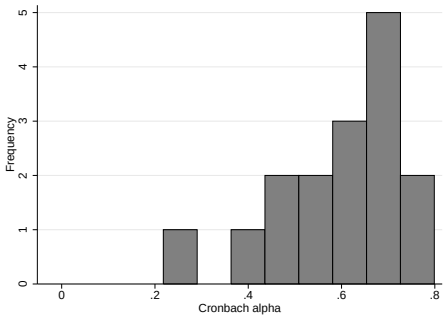

(d) Question 3

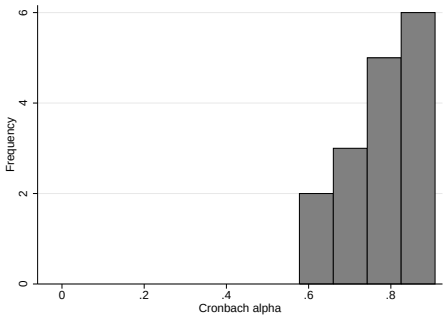

Figure 14: Cronbach alpha's for the 161 vacancy pairs

(a) All questions

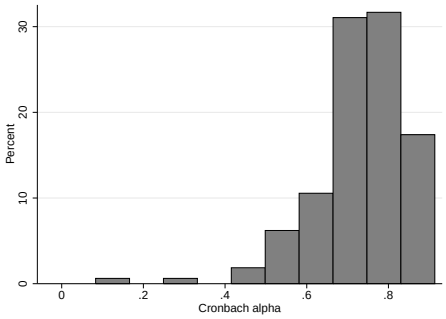

(b) Question 1

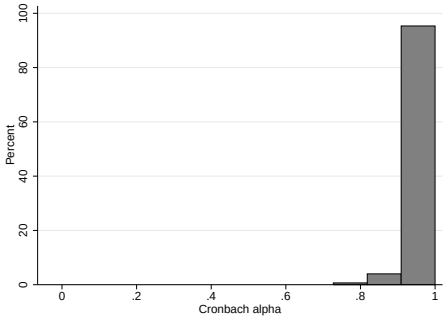

(c) Question 2

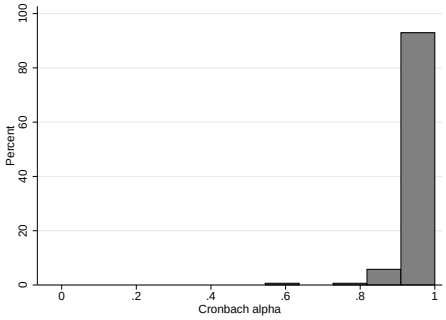

(d) Question 3

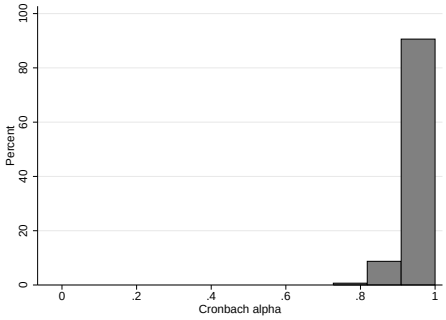

Supplemental Online Appendix - Instructions and Consent Forms

OA.1 Consent form

Consent Form for Participants: "How Do Unemployed Search for Jobs?"

Thank you for your willingness to consider taking part in this study. Please read the information below carefully. By signing the consent form below, you indicate that you have understood the purpose of the study, you have been made aware of your rights and you have agreed with the terms and conditions of the study.

Purpose of the study

The study is undertaken to understand better how people search for jobs. The study aims to observe how people search for real jobs. The goal is to document parts of the job search process.

How will this work?

The study will be conducted over a period of 12 weeks and you are asked to take part to one weekly session of 2 hours taking place at a pre-agreed time slot. You will be asked to come to our computer facilities, located at the School of Economics, 31 Buccleuch Place, EH8 9JT Edinburgh. There will be a maximum of 30 participants present at the same time in the facilities. The research team aims to provide an environment that is conducive to the job search of participants and hopes that participants will attend for the duration of the study or up to the point you find a job.

You will be able to spend most time each week to search for job vacancies. These job vacancies are obtained from two sources:

- Our main data source is the vacancy database of Universal Jobmatch and coincides with those used at Jobcentre Plus.
- Additionally, our database includes a small number of vacancies (no more than 2 per 100 vacancies) that is added for research purposes. These "research vacancies" are included to understand better which types of vacancies people are interested in even if these are not currently offered. If you express interest in such a vacancy, you will be immediately informed that this is a research vacancy before you start any application.

We will track the pages you consult, what vacancies you are looking at and consider applying to. This information will never be linked to any of your personal information such as your name and address, which will be stored separately. Your personal information will never be given out to anyone and will be accessible only to selected members of the research team.

You will also be asked some survey questions about your job search in the past week and your wellbeing. In the initial week, we will also ask a number of questions about your background and unemployment history. Six month after the end of your participation we will send you a survey about your labour market experience and your well-being.

Note that we ask all participants to stay for the full 2 hours in the laboratory. But if you do not want to search for jobs anymore, we provide some alternative ways in which you can use the computer and internet facilities.

If you are unable to participate to a session, please inform us as soon as possible (under jobsearch@ed.ac.uk or 0131 6508324). The research team will attempt to provide additional slots in case a participant misses his time slots for justified reasons (e.g., job interviews, illness).

Important notes

- Participation to this study is entirely voluntary. You should by no means feel complied to participate. You can also withdraw from the study at any time if you wish to do so.

- Since the study is to gain understanding in how people search for jobs, the research team holds no particular view on how individuals should search for jobs. Thus, you should search for jobs in the same way as you would normally do.

- The study is conducted by the research team, and no personalized information is shared with any other organization. Therefore, no information will be shared with Job Centre Plus or the Department of Work and Pensions. If you would like to obtain a record of your search activities, e.g. to use for discussion with your case worker, you can obtain a printed record to take along at the end of each session.

- You should be aware that **participation in this study does not provide any additional benefits**, and in particular it does not provide particular help in job search. In particular, you **should follow your usual job search strategy**, such as for example looking at other job vacancies beyond those provided in our database, searching from home via the internet, and contacting friends and acquaintances. You should not take the time within the study as an indication of the appropriate time to spend on searching for a job.

- All the data collected during your time in our computer facility is anonymous. Your search activities will not be matched to your identity in any way. You will be attributed a randomly generated number at the first session and all data records will be matched to that number.

- We will ask you for a telephone number that we can use to contact you. We will only contact you to remind you of the time slot you have been allocated to and to inform you of any changes in schedule. Of course the telephone number will not be matched to the data we collect in the laboratory.

- You have the right to withdraw entirely from the study (i.e. ask us to delete all the data records associated with you) at any point during the study.

- The impersonal data collected will be used for research purposes (and ONLY for research purposes). Personal data will never be given out, and will be eliminated after the study is completed. The results of the study will be published in peer-reviewed scientific journals.

Compensation

You will be compensated for your efforts of coming to and participating in each session in our computer facility with a compensation of £12.50 per visit (2 hours) to the laboratory. Additionally, if you participated in all four sessions in the first four weeks you are entitled to a £50 clothing voucher for job market attire as compensation for arranging the visit every week. The same holds for weeks 5 to 8 and for weeks 9 to 12.

Eligibility

Participants have to be at least 18 years of age, permanent residents of the UK and living in Edinburgh (or within a distance of 5 miles from Edinburgh). You should be seeking for a job for a period of 4 weeks or less at the start date of the study.

Signature

If any of the material above is unclear to you, or if you have any doubts and would like clarification, please consult a member of the research team before proceeding.

If you are willing to take part in this study, please sign the consent form below:

I certify that I voluntarily participate in this research study. I certify that I read and understood the information above, and am eligible for taking part in this study.

(please print your name)

(please sign)

(place and time of signature)

OA.2 Lab instructions

UNIVERSITY JOB SEARCH STUDY: INSTRUCTIONS

Please do not start using the computer before we indicate you to do so.

We will read these instructions aloud at the start of the first session.

INTRODUCTION

Welcome and thank you for coming here today. Before we explain how each session will work, we would like to raise your attention to the following:

- **Health and Safety**: There will always be one person from the research team in the computer room. There is one toilet on this floor that you are free to use. In case of fire, please do follow the signs for fire exit. The main exit is through the staircase you have used to come up here.
- **No smoking**: Smoking is not allowed in this building.
- **Silence**: Since there are many of you in the room, we would appreciate if you would keep silent, so that everyone can concentrate on their computer activity.
- **Mobile phones**: Mobile phones must either be switched off or be on "silent" during each session. We would appreciate if you leave it on only if you are expecting an important phone call. And if you do receive a phone call, please leave the room and take the call outside (in the staircase).
- **Food and drinks** are not allowed in this room.
- **Questions**: Please do not hesitate to call us if you have a question.

WHAT IS THE STUDY ABOUT?

The goal of the study is to understand how people search for jobs. Importantly, we hold no preconceptions regarding how people *should* search for jobs. We designed this study to find out what people usually do and what strategies are most successful. At the moment, we do not know what these are. We are interested in finding out common patterns in search strategies, and kindly ask you to search exactly in the same way as you normally would.

WHAT WILL HAPPEN IN EACH SESSION

When you come in, you will be assigned to a computer station. We may provide specific instructions at the beginning of the session, so please do wait for us to indicate the start of the session. We will now describe how each session will proceed.

1. LOGIN

You have received a unique login number and password that you can use to login on the website here and also from home. You will be able to access your records using this login information.

2. SURVEY

Each weekly session will start with a **short survey**, asking questions about your past week and job search. After filling the survey, you will be re-directed towards the job search engine's main page.

For the first session, we will ask you to fill in a longer survey asking you questions about your background, qualifications and job search experience so far. You will only need to answer this initial survey once, in this session. It should take 20 minutes to fill in this initial survey.

3. THE JOB SEARCH ENGINE

We have designed our own job search engine. It allows you to search through all UK vacancies that are also recorded in Universal Jobmatch.

We ask you to search for jobs using this search engine only for a minimum of 30 minutes.

You can search using various criteria (keywords, occupations, location, salary, preferred hours). Importantly, you do not have to specify all of these. You just need to fill at least one of them.

If you specify more than one criterion, it is important to note that the computer will search for vacancies that satisfy all the criteria at the same time. For example, if you enter a keyword and you also select an occupation, it will search for vacancies that match both at the same time. Vacancies that match the keyword but not the occupation will not be shown.

Within some categories you can fill in more than one field. For example, within "occupations" you can specify up to two of them. If you do fill in two occupations, the computer that match either the first OR the second occupation. Vacancies that match one occupation but not the other will still be shown. You can also specify more than one pay range. This allows you to specify, for example, the hourly wages and the yearly wages that you are willing to accept. If you only specify hourly wages, it will not show vacancies that only specify yearly wages.

If you fill in your preferred hours, for example full time work, it will only list vacancies where the employer ticked a box that it is full-time work. Vacancies where the employer did not explicitly state that it is full-time work will not be shown.

If you leave a field empty, the computer will not use that criterion to restrict your search.

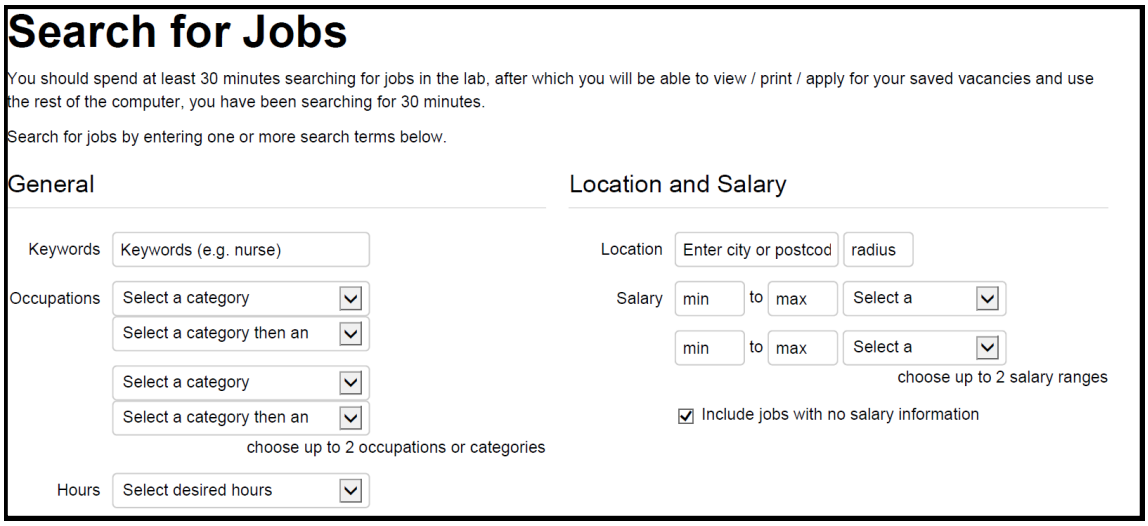

Once you have defined your search criteria, you can press the search button at the bottom of the screen and a list of vacancies fitting your criteria will appear. You can click on each individual vacancy to get more information about it. You can then either

- **Save the job (if you are interested in applying)**
- **Do not save the job (if you are not interested)**

If you save the job, the computer will keep a record of the vacancy. You will be able to see all records of all saved vacancies at the end of the session.

If you do not want to save the job and want to go back to the search results, we will first ask you a few questions about why you are not interested in the job. Your answers are very important to us.

You can modify your search criteria at any point and launch a new search.

Note that we have also created a small number of vacancies ourselves (about 2% of the database), which are there for research purposes only. This is to learn whether you would find these vacancies attractive and would consider applying to them if they were available. We kept them to a minimum not to disturb your search. These vacancies will appear as all the other vacancies and may appear in your search results. But we will inform you at the end of the 30 minutes of any vacancy that may not be real. You will be able to see the list of your saved vacancies immediately after the 30 minutes are over, and we will indicate if any of them was an artificial one.

We may try alternative interfaces for the job search engine in the coming weeks. We will inform you if we do so and will explain the changes at that point in time.

4. FREE USE OF THE FACILITIES (after 30 minutes)

We will let you know when the first 30 minutes are over. You will then be free to use the computer for other purposes. You can of course keep searching using our job search engine, or you can do other things, such as write your CV, write a letter, or even send e-mails. You can use the facilities for up to 2 hours.

If you do not wish to continue searching or use the computer for other purposes, you are free to leave.

END OF THE SESSION

We can print a record of your job search for the day (just call us once you have finished), but only if that is your wish. You are free to show these records to your adviser at the Job Centre. They informed us that this would count as a proof of search activity.

Compensation: In general, you will receive a total of £11 as a compensation for your travel and meal expenses. This time, as you will soon discover in the initial survey, we do offer you the possibility of investing part of this compensation in this initial session. This is not compulsory. But if you do choose an investment option, your earnings will then be a function of what investment you have chosen.

Please collect your compensation from the registration room. You will get an envelope and be asked to sign a receipt. Note that the Job Centre has agreed that these £11 are a compensation for expenses and are not an income.

IMPORTANT NOTES

LOG IN FROM HOME OR FROM ANOTHER COMPUTER

You will be able to use our search engine from home or from another computer as well. You just need to log in on the website and use your login information. You will be able to see all the vacancies you saved and will be able to retrieve all the relevant information about them.

Note that as indicated in the consent form, all records saved are anonymous. These will not be matched to your names at any point.

YOUR COMMITMENT

Note that it is very important for us that you come back every week and search in our facilities, unless of course you have found a job. If for one reason or the other you do have to cancel your session in a given week, please let us know as soon as possible. We will either try to reallocate you to another slot or ask you to search from home in that particular week. If you have found a job, please do let us know. This is of course of key importance for our study.

Also, importantly, you will receive a £50 clothing voucher for each four consecutive weeks you come. The first voucher will be distributed in the fourth week, that is, three weeks from now. The second voucher will be distributed in the eighth week and the third voucher in the twelfth week.

Thank you very much for your attention. If you have any questions, please raise your hand and we will come to you.