

DISCUSSION PAPER SERIES

IZA DP No. 11814

**How Wage Announcements Affect Job
Search: A Field Experiment**

Michele Belot
Philipp Kircher
Paul Muller

SEPTEMBER 2018

DISCUSSION PAPER SERIES

IZA DP No. 11814

How Wage Announcements Affect Job Search: A Field Experiment

Michele Belot

University of Edinburgh, European University Institute and IZA

Philipp Kircher

University of Edinburgh and European University Institute

Paul Muller

University of Gothenburg

SEPTEMBER 2018

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

How Wage Announcements Affect Job Search: A Field Experiment*

We study how job seekers respond to wage announcements by assigning wages randomly to pairs of otherwise similar vacancies in a large number of professions. High wage vacancies attract more interest, in contrast with much of the evidence based on observational data. Some applicants only show interest in the low wage vacancy even when they were exposed to both. Both findings are core predictions of theories of directed/competitive search where workers trade off the wage with the perceived competition for the job. A calibrated model with multiple applications and on-the-job search induces magnitudes broadly in line with the empirical findings.

JEL Classification: J31, J63, J64, C93

Keywords: online job search, directed search, wage competition, field experiments

Corresponding author:

Michele Belot
School of Economics
University of Edinburgh
30 Buccleuch Place
Edinburgh EH8 9JT
United Kingdom
E-mail: michele.belot@ed.ac.uk

* 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. We benefitted from helpful comments of many seminar audiences including Princeton, Minneapolis FED, IFS, CRAI and University of Copenhagen. Kircher acknowledges the generous financial support from the European Research Council Grant No. 284119, without which this study would not have been possible.

”Though wages in bargaining models are completely flexible, these wages have nonetheless been denuded of any allocative or signaling function: this is because matching takes place before bargaining and so search effectively precedes wage-setting. [...] In conventional market situations, by contrast, firms design their wage offers in competition with other firms to profitably attract employees; that is, wage setting occurs prior to search so that firms can influence the allocation of resources in the market.” Hosios, 1990, p. 280.

1 Introduction

Wages represent the price of labor and are key to the efficient functioning of the labor market. Nevertheless, we know relatively little about how wages affect the matching process between vacant firms and unemployed workers in the private sector even at a very basic level. From the firms’ perspective, would an increase of the wage for a given job posting increase the number of applicants? From the workers’ perspective, do some of them apply only to lower-wage firms (presumably as a strategy to avoid competition at high wage firms)? We conduct an audit study in a mid-scale field experimental setting where we randomly assign wages across jobs posting in a large number of occupations, and provide affirmative evidence.

The two simple questions we investigate provide important information about the workings of the labor market. How firms benefit from wage offers and how workers react to them are key to understand incentives and efficiency in such markets. The literature on directed search - also called competitive search¹ - grew out of a presumption that wages may play a key role in the functioning of the labor market. In its elementary form, such theories posit two features. Firms will strategically increase the wage to increase the expected number of suitable job applicants up to the point where the additional wage bill becomes too large. The underlying mechanisms why applications increase with the wage is that workers strategically direct their applications more often to higher wages. They do so up to the point where the gain from a higher wage is offset by the additional competition - and this implies that not all workers apply to the top wage as competition would be too tough there while it would be too weak at slightly lower wages. The last argument relies on the key assumption in the directed search literature that time or resource constraints prevent workers from applying to all possible jobs, implying selectivity in where they apply.

In his seminal contribution on efficiency in search markets, Hosios (1990) was the first to emphasize these directed search features as central for efficiency, following his main observation that in general efficiency fails under random search in the Diamond-Mortensen-Pissarides tradition where these features are absent, necessitating many types of policy interventions.² His quote at the beginning of this article asserts that ”conventional markets” use the wage to attract applicants, and the truth of this

¹For an extensive survey of the competitive/directed search literature see Wright et al. (2017).

²Hosios (1990) is mostly known for its (in)efficiency result under random search, where he also briefly touches on policy interventions (this is a large theme in the subsequent literature, see e.g., Mortensen and Pissarides (2003). Less well known is Hosios’ study of directed search (pp 293-296) and his confirmation about the efficiency of this model class. The efficiency-enhancing role of competitive/directed search has been subsequently stressed, amongst others, by Peters (1997), Moen (1997), Acemoglu and Shimer (1999), Shi (2001), Mortensen and Wright (2002), Kircher (2009), Menzio and Shi (2011), Gourio and Rudanko (2014), or Lester et al. (2015).

assertion constitutes part of our research question, together with some exploration of the mechanisms that underlie it.

Despite the simplicity and importance of the questions, it is surprisingly difficult to obtain answers from observational data because variation in wages tends to be correlated with variation in other features of the job announcements, as we discuss below. To overcome this, we adapt the "audit study" methodology that has recently seen a resurgence to understand how employers distinguish between workers that only differ in a single dimension of interest.³ Here we apply it to the other side of the market to understand how unemployed job seekers distinguish between private sector jobs to which we randomly assign different wages. These artificial vacancies are sown within a large set of real job announcements. We conduct the study on a job-matching platform that we set up for this very purpose. We find a positive answer to both questions. Not only the qualitative finding but also the economic magnitude is broadly in line with a calibrated directed search model in which workers adjust their search behavior to the wage and the perceived competition from other workers. We also ask another sample of non-student participants to rate job announcements, and higher wage jobs are indeed perceived as more competitive, which is the key trade-off for workers in the model.

The audit study approach to assign the wage randomly to artificial vacancies does raise ethical issues, particularly because the population of job seekers is a vulnerable population. Actually creating jobs corresponding to the posted vacancies, as in Dal Bo et al (2013, reviewed below), is feasible for a specific job (like in their case particular work with disadvantaged children for the Mexican government), but again remains infeasible across a large number of private sector occupations. Our Ethical Review Board would not authorize a study design without prior consent from the job seekers (this is in contrast with audit studies on the employer side that are conducted without consent). For this reason, we opted for a medium scale field experiment.

We created our own job matching platform, providing access to up-to-date vacancies that were kindly shared with us by Universal Jobmatch, the UK government's job search site which holds one of the most comprehensive vacancy datasets in the UK. We then recruited 300 recently unemployed job seekers to participate in our study at the experimental laboratory in Edinburgh once a week for a duration of 12 weeks. Per session they spent 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. Although being logistically demanding, creating our own job platform and running the study in our lab has significant advantages though. Most importantly, we are able to record precise information on job search in a way that is tailored to our study. On top of that, we can ask participants explicitly for consent, verify their identities and distribute compensation for participation.

As one element in an extensive initial introduction, participants were informed that a small fraction (less than 2% of vacancies) would be posted for research purposes to understand whether they would be interested in these jobs if they were actually available. This percentage is intentionally small as

³Traditional audit studies trained pairs of individuals to act as similar as possible and to differ in only one dimension such as race or gender (see for example Neumark et al. (1996), who consider gender discrimination in the labor market and Ayres and Siegelman (1995) who apply the same approach to a car sales setting). More recently the focus has shifted to sending out fictional resumes. Bertrand and Mullainathan (2004) consider ethnic discrimination in the hiring process by varying the names of the applicants between 'African-American'- and 'White-American'-sounding. Many other aspects have been studied since, including age (Lahey (2008)), gender (Petit (2007)), physical appearance (Rooth (2009)), and unemployment duration (Kroft et al. (2013)).

not to affect individual’s job search incentives, but high enough to generate enough vacancies for a meaningful analysis. For ethical reasons we inform individuals after the search phase but before they actually apply whether a saved vacancy was posted by us for research purposes. It implies that we do not observe actual applications, and rather study saving (or viewing) a vacancy, which for actual vacancies is a strong indicator of interest in the vacancy and likelihood of applying.⁴ In this respect, our study is similar to audit studies for resumes, where the outcome is the callback rate rather than actual interviews or job offers.

Job seekers were informed about the source of our regular vacancies. We did not specify the exact nature of our research vacancies. In praxis, these were expired vacancies from at least half a year prior to our study that did not allow direct firm identification and for which we artificially multiplied the original wage by some factor. We created pairs of vacancies that are nearly identical except for minor changes in wording and are posted at roughly the same time, and assigned either the original or the multiplied wage at random. This allows us to answer the first research question by assessing the increase in applications with the wage within a pair. It also allows answering the second research question whether there exist people who apply only to the low wage even if higher wage jobs are present. This is a key feature of directed search, while reservation strategies whereby individuals aim for any job above a cutoff wage are common in random search.⁵

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 0.7% - 0.9% more saves. Second, we find that 42 % of those that save the low wage vacancy within a vacancy-pair do not also save the high wage vacancy. This percentage remains almost constant when conditioning on individuals that were shown both vacancies on their screen (39%). We also find that more recent vacancies attract a lot more interest, even if they are posted only a day apart. Controlling for this makes little difference for the wage elasticity of applications. Moreover, even if the high wage vacancy is posted more recently, still 16% of individuals who end up saving the low wage express no interest in the high wage one. When using viewing of the vacancy (rather than saving) as the outcome the results are very similar. In a robustness analysis we exclude that this finding is due to (1) study participants identifying the research vacancies, (2) differences in location between the vacancies in a pair or (3) learning over time. We note that our findings on the wage elasticity are very robust across specifications. On the other hand, the magnitude of the share that saves the low wage vacancy but not the high wage vacancy varies quite a bit across specifications (though it is always significantly larger than zero).

Our randomized experimental setup ensures that both vacancies within a pair are observationally virtually identical in terms of all other aspects of the job. But of course it is plausible that participants *perceive* these vacancies differently, that is, that they use the wage as a signal of important aspects of

⁴In our study, almost one-third of all saved (real) vacancies is eventually applied to.

⁵Stationary random search models where jobs offers come with different take-it-or-leave-it wage offers in the spirit of McCall (1962) or Burdett and Mortensen (1998) yield a reservation strategy: If a person accepts a low wage job s/he also accepts the same job at a higher wage. This usually concerns the *acceptance* stage of a job offer, so applications may still be random. But we were concerned that individuals with a tiny cost of applying for jobs will only do so if the wage offer is above their reservation threshold. This would yield a reservation strategy in terms of their job application decision. In contrast, in the canonical directed search model workers have only one job application and therefore obviously someone who applies to a low wage job has to forgo applying to a higher wage job with probability one. This selectivity persists but at a less stark level when workers can send multiple applications.

the vacancy. Directed search theories suggest that high wage vacancies should attract more interest (and from better workers), and therefore high wage vacancies could be perceived as such. There could however be other signals attached to high wage vacancies, such as worse working conditions, in line with a compensating differentials hypothesis. To understand better how vacancies are perceived, we designed and conducted a complementary survey (with different participants). We find that the high wage vacancy within a pair is 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) goes against 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 as in Albrecht et al. (2006) and on-the-job search as in Postel-Vinay and Robin (2002). Both are important to allow for sensible magnitudes in the application behavior. Only if workers can apply to multiple jobs is it possible that - conditional on applying to a low wage job - at least some might also apply to a high wage job. If workers can search for additional offers on the job the initial wage difference between any two jobs is less important to their life-time wealth, inducing workers to queue at substantially different initial wages (otherwise the model cannot sustain substantial wage dispersion for reasons akin to those pointed out by Hornstein et al. (2011)). To our knowledge we are the first to combine both features in a directed search model, and despite this the model remains tractable. 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 elasticity with respect to an out-of-equilibrium wage change is very close to what we find in our study, and agents in the model do not apply to high wage vacancies conditional on applying to a low wage vacancy to a large degree because of the competition that they anticipate.

Our model does not speak to the interesting finding that job seekers save more recent vacancies much more, even if both are new to them and both were listed on their screen. The reason is a lack of models to build on for this new finding. The only work we are aware of explains this by the fear of job seekers that older jobs might have already been filled in. Albrecht et al. (2017) build a competitive search model with this feature, and their calibration suggests a rapid decline in job applications between the first and second day that exceeds even our sizable measure of 37% decline over that period. This indicates that competitive search could be fruitful to understand this novel phenomenon, yet marrying their approach with multiple job applications and on-the-job search exceeds the confines of this - mostly empirical - paper.

The rest of the paper is organized as follows. Section 2 uses the deeper comparison with the related empirical literature to discuss further benefits and limitations of our approach. The subsequent section explains the experimental setup in general and how the artificial vacancies were created and posted in particular. Section 4 presents the main empirical analysis and results as well as a large number of robustness checks. In Section 5 a simple directed search model is laid out and its predictions are compared to our empirical findings. Section 6 concludes with a summary and an outlook how this

type of study could be used for wider questions.

2 Discussion of our approach relative to the literature

As mentioned above, the idea to investigate whether a firm that offers a higher wage will attract more job applicants is not new, even though the number of studies is limited as there are only few datasets that include application behavior or other measures of interest in a job.⁶ Existing studies using observational data usually observe different vacancies with different wages. The key challenge is to ensure that these are identical except for the wage, as they would be if an employer changed the wage on a given vacancy. Otherwise one obtains counter-intuitive outcomes: Faberman and Menzio (2016) exploit a rich survey performed in 1980-82 and find a negative relationship between the starting wage of a vacancy and the number of applicants, even after controlling for three-digit occupations. This is in line with early results by Holzer et al. (1991) who find that firms that pay the minimum wage receive more applicants than firms that pay slightly more. Marinescu and Wolthoff (2014) replicate such a negative relationship on a much larger dataset from Careerbuilder.com, again controlling for occupational codes. Yet after controlling for the much more detailed job description, the relationship reverses: higher wage attract more applicants. Even job titles do not seem to be sufficient to make job ads comparable: Banfi and Villena-Roldán (2016) show for Chilean data from *trabajando.com* that intended wage payments positively correlate with the number of job applications after controlling for many observables including job title, even in cases where those wages are not actually shown to the job seekers. In some sense that is encouraging because firms evidently are able to communicate the attractiveness of their job even in the absence of explicit wage posting, but it also highlights the difficulty of making observational vacancies sufficiently comparable as job seekers seem able to distinguish jobs even with identical job title. Therefore, we propose an alternative strategy here.

Conducting this audit study as a field experiment has several advantages on top of the main benefit of ensuring orthogonality between the job description and the wage. Because search is carried out on our own job search site we observe actual search behavior, which allows us to assess whether an individual who chooses only a low wage vacancy did actually encounter the twin "higher-wage" vacancy on his screen. Second, it allows us to ask for consent to ensure ethical approval, which is challenging when a study concerns posting artificial vacancies. Third, the fact that we consider artificial vacancies implies that we can post vacancies in a wide range of different occupations and skill levels, rather than focusing on one narrowly defined job type.

This distinguishes our approach from studies that actually provide the underlying jobs but are therefore restricted to specific occupations. The influential study by Dal Bó et al. (2013) considers mission-motivation for civil servants working with disadvantaged children in rural Mexico. They find that higher wages increase the number of attendees in an assessment centre, their quality and motivation. Interestingly, the wage elasticity of around 0.7 that we find across a large number of occupations in the UK is similar to their finding for this specific occupation in Mexico, and is also similar to the non-experimental elasticity reported in Marinescu and Wolthoff (2014) for the US once they control for job title. Abebe et al. (2017) advertise three-months clerical positions in Ethiopia and find qualita-

⁶A different theoretical approach to test between random vs competitive search based on job-to-job transitions is proposed in Lentz and Moen (2017), though they have not yet applied it to actual data.

tively similar results to Dal Bó et al. (2013), though the wage elasticity of assessment centre attendance is substantially lower, but in their case the job is temporary and one might conjecture that a similar salary increase for a permanent job might lead to a higher response. They also document particularly high application costs in their setting and show benefits of application subsidies.

In a different context and with a different research focus, Leibbrandt and List (2015) post (real) vacancies for administrative assistant as male or female jobs and vary whether they explicitly mention that wages are negotiable (but they do not mention any salary level). They find differences by gender on applications and negotiation behavior. Mas and Pallais (2016) vary the relative wage for flexible working hours for call centre jobs, and find limited willingness to pay for flexibility. None of these studies varies the wage in the job announcements, but rather at an interim stage after individuals have already contacted the employer, and so findings could be consistent even with random search. None of them posts pairs of jobs with different wages at the same time to see how the presence of a high wage offer affects the behavior towards the low wage, and vice versa.⁷

In contrast to the existing studies, our setting involves many occupations but the announcements are not real. We chose a small number of job postings to not alter the usual search patterns of job seekers, and we have no reason to believe that it did alter it. First, it is not the case that our intervention spoiled an otherwise pristine set of real job vacancies. Rather, the presence of fake job advertisements are routinely reported for websites like Monster.com, Careerbuilder.com and Universal Jobmatch, and advice for job search on the internet usually cautions about this. For our particular vacancy source from Universal Jobmatch, investigative journalism in the UK places the lower bound of non-real vacancies at 2% even in the absence of our intervention, while the tabloid press warns of even higher numbers.⁸ Nevertheless, the database is regarded as reliable by the UK government, who encourages job seekers to use it and repeatedly contemplated to make search on this platform mandatory for recipients of Job Seeker Allowance. We take from this that a low number of non-real vacancies are standard in current online job search. Moreover, in an exit survey the vast majority of our participants report that the presence of research vacancies was immaterial for how they search for jobs, and that they were unable to distinguish them from real ones.

Obviously the restriction that we could not conduct our study in a regular market setting implies limitations in sample size. Therefore, for this study we evaluate average effects and refrain more detailed questions relating to heterogeneous job search behaviour along different dimensions which require more power. We hope this will become feasible in future studies. If larger future studies were approved, the general methodology we propose here could also be used to study wider questions, such as whether minorities react differently to higher wages or employer ethnicity or gender.

3 Experimental Design

The setup, recruitment process and institutional setting are in common with our other paper Belot et al. (2016), which focuses on the overall job search behavior of participants within this study and

⁷There are also studies on the labor supply of existing workers rather than new hires: for example, Fehr and Goette (2007) and Goldberg (2016) vary the hourly wage rate for bicycle messengers in Zurich and for agricultural workers in Malawi, respectively. They study their subsequent choice of hours.

⁸ See for example Channel 4 (2014) and Computer Business Review (2014).

on the role of providing tailored advice. The experimental intervention evaluated in this other paper is orthogonal to the one studied in the current paper. We reproduce here the relevant aspects of the design. Further details of our setup can be found in the other paper.

3.1 Real Vacancies and Artificial Vacancies

3.1.1 Real Vacancies

In order to provide a realistic job search environment, we created a job search engine that accesses a local copy of the database of real job vacancies of the government website Universal Jobmatch. This is one of 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. Each week between 800 and 1600 new vacancies were posted in Edinburgh. Comparing our database with UK national vacancy statistics suggests that it contains over 80 % of UK vacancies.⁹ This is a very extensive coverage compared to other online platforms. 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.

3.1.2 Artificial vacancies


A small number of artificial vacancies was introduced during the study. Participants were fully informed about this. While the main introductory message to participants truthfully conveyed our interest in studying how people search for jobs - covered in our companion paper Belot et al. (2016) - they were also 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.¹⁰

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. Each artificial vacancy was only active during sessions of a particular week, such that participants would never observe them in multiple sessions. In this section we describe the

⁹Based on 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

¹⁰In 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 on online job search sites (see 8 and in this connection Craigslist’s chief executive Jim Buckmaster is quoted: “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 (2009)). This worry is routinely mentioned to job seekers in many search guidelines (see e.g. Joyce (2015)).

Figure 1: Occupational distribution of real and artificial vacancies




procedure used to create the artificial vacancies and present some statistics on comparability to the set of real vacancies.

3.1.3 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 description of the vacancy had 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 all identifying information (company name, contact person, telephone number, website, etc.).¹¹ Since such information is often missing in vacancies in our sample, we do not believe that we moved out of the ordinary with this. 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

¹¹We 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.

Figure 2: Distribution of posted salaries



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 42% of all vacancies on Universal Jobmatch post a salary, and vacancies that post salaries may differ from those that do not. Figure 1 shows the distribution of vacancies across occupations. The left (light) bars denote all real vacancies that were posted on Universal Jobmatch during the study.¹² 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 middle bars of Figure 1. The distribution of vacancies with posted salaries is quite similar to the overall distribution, with only occupations group 5 (Skilled trade) being less likely to post a salary. The second step in the selection procedure required vacancies to have a ‘simple’ description that allows easy manipulation to ensure anonymity of the employer. To select suitable vacancies, we went through a set of outdated vacancies posted on Universal Jobmatch, and checked one by one whether a vacancy was simple enough to manipulate. Clearly, the vacancies that we selected are not representative of all vacancies posted on Universal Jobmatch. We show the occupational distribution of all artificial vacancies in right (darkest) bars of Figure 1.¹³ We oversample vacancies from occupational group 5 (Skilled Trades) and from groups 9 (Elementary Occupations). This is not surprising as jobs in these categories typically have a shorter description, making them easier to manipulate. Still, there is considerable variation across occupations, as the majority of vacancies are still posted in the other occupations and all occupations remain represented.

Employers can also choose to post a single “point wage” or they can post a “wage range” from a

¹²It is based on a sample of 30,000 vacancies posted in the Edinburgh area around the start of the study.

¹³Note that the SOC code of the vacancy is not always ‘correctly’ specified by the employer. To keep the artificial vacancies as close as possible to the real vacancies, we did not correct the codes. Here we want to provide an accurate overview of the occupational distribution of the vacancies and therefore we have ‘corrected’ the codes for those vacancies that seemed to be incorrectly classified. This correction has been performed using the Computer Assisted Structured Coding Tool (‘CASCOT’), provided by the Warwick Institute for Employment Research. Results for the raw coding are similar, but suggest a better representation in high-skilled occupations.

minimum to a maximum. We report consistently the minimum wage, which coincides with the point wage if only one wage is offered. 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 Figure 2. Panel (a) compares annual salaries of real vacancies and artificial vacancies and panel (b) does the same for hourly salaries. The artificial salaries are those *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.¹⁴

3.1.4 Manipulation of wages and locations

Our strategy is to create random variation in the posted salary within the chosen set of vacancies while keeping all other vacancy characteristics constant. We create pairs of vacancies. Both share all key vacancy attributes except for the posted wage.

This approach is parallel to the randomized audit studies in which pairs of applicant’s résumés 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 resumes from a pair. The other resume is sent to a different employer. We make both vacancies accessible here because it allows us to directly test the hypothesis that job seekers strategically use the wage information to target only one of otherwise similar vacancies. Specifically, we will be able to see whether some job seekers only consider the low wage vacancy even though they have also been exposed to the high wage one. Of course, the use of pairs of artificial vacancies also allows to filter out unobserved characteristics of the vacancy, which improves the precision of the wage elasticity. 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. See the online appendix OA.4 for two examples of vacancy pairs. 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.

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%.^{15,16} The wage assignment was conducted in two stages. First it was randomly decided which vacancy would have a changed salary, second it was 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.¹⁷ As a result our sample contains relatively more wage increases among low wage vacancies than among high wage vacancies, while also overall


¹⁴Only 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.

¹⁵Note that this is in the same order of magnitude as the wage increase implemented by Dal Bó et al. (2013), which is 33 %.

¹⁶One may worry that adding or subtracting a percentage leads to unrounded numbers that might look suspicious to a job seeker. This is not the case though, since the original wages of the vacancies that we created were not rounded numbers in general either.

¹⁷In case the assigned wage decrease resulted in a wage below the minimum wage, we assigned a (random) wage increase instead. We did so to prevent the vacancy from looking suspicious, though the set of real vacancies actually contains posted wages below the minimum wage.

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



we have more wage increases than decreases. Two waves of job seekers were confronted with these vacancies. During the second wave of the study, the same set of artificial vacancies was used, however the wage was switched around within the pair. In total, we created 322 vacancies (161 pairs), based on 94 original vacancies.¹⁸ The wage was reduced in 32 pairs, it was increased with 20% in 75 pairs and increased with 40% in 54 pairs.¹⁹ For vacancies that originally offered a wage range, we multiplied both their minimum and maximum wage by the same factor.

While within the theory of competitive search it is common to study the reaction to wages that lie off the equilibrium path, for empirical purposes it is important that the wage manipulation stays within a reasonable range, so that both vacancies appear realistic. To show that this is the case, we present a measure of wage dispersion at the finest (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 within this occupation.²⁰ We find that on average a 20% wage increase or decrease corresponds to 0.44 of a standard deviation, while a 40% wage change corresponds to 0.88 of a standard deviation. The distribution of these numbers across occupations is shown in Figure 3. 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

¹⁸In addition, we 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 these pairs and we focus our analysis on the Edinburgh pairs. Furthermore, 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 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 20, 40 or 60%. All results in this paper are, unless mentioned otherwise, based on “Edinburgh pair” artificial vacancies only. In the appendix we show that including the Glasgow pairs does not change our results.

¹⁹Initially 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.

²⁰These computations are based on vacancies that post a minimum annual wage above £1000, and 9 occupations (out of 72) for which we have less than 5 observations are excluded.

artificial vacancies from real vacancies, 68% answered ‘never’ or ‘rarely’. We confirm robustness of our results when restricted to only this sample. 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 4.3.1 we show that there is also no indication of any learning over the duration of our study among job seekers in terms of identifying artificial vacancies.

3.2 Recruitment of Participants

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 was announced but 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 were recruited at local public unemployment agencies (Job Centres) and received unemployment benefits (Job Seekers Allowance, JSA). Since JSA claimants tend to see their advisers every two weeks and we recruited for two weeks in the Job Centres, it is likely that we approached most of them. Amongst eligible individuals that we approached, about half signed up and of these half actually appeared for the study.

In Table 1 we present characteristics of our participants collected at baseline in the first week of the study. The top panel displays demographics and the middle panel displays summary statistics of their job search history. Whenever available, we compare them to average characteristics of the population of job seekers in Edinburgh. The population statistics are retrieved from the NOMIS database of JSA claimants. We focus on those with unemployment duration up to 6 months, because for these the median unemployment duration is almost equal to that of our participants (80 days). Only a limited number of characteristics is available for this group in NOMIS. Our study slightly oversamples females and non-whites, while the average age is very close to the population average. We have a fair representation of participants with or without higher education, but lack a comparable statistic in the population. In terms of job search history, participants indicate to have applied to 64 jobs on average and have attended 0.52 interviews. In the lower panel we show some summary statistics for the job search behavior that we observe during the study. Per week, participants *list* on average 528 vacancies which are displayed on their screen, of which they *view* 24 in detail, *save* 9.9, *apply* to 2.4 and obtain 0.076 job interviews. We explain these search activities in more detail below. In addition they search through other channels besides our website (resulting in 7.7 applications per week).

3.3 Job Search

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.

Table 1: Characteristics of study participants

	Study participants				Population ^a
	mean	sd	min	max	mean
Job search history: ^b					
expect job within 12 weeks (%)	58	49			
vacancies applied for	64	140	0	1000	
interviews attended	0.52	0.91	0	10	
jobs offered	0.42	1.1	0	8	
days unemployed (mean)	260	620	1	5141	111
days unemployed (median)	80				81
Average weekly search activities:					
listed	528	449	3	3968	
viewed	24	16	2.5	119	
saved	9.9	10	0	92	
applications (in lab)	2.4	4.1	0	37	
interviews (in lab)	0.076	0.24	0	2.8	
applications other	7.7	8.6	0	50	
interviews other	0.51	0.86	0	11	
Demographics:					
female (%)	43	50			33
age	36	12	18	64	35
high educ ^c (%)	43	50			
white (%)	80	40			89
couple (%)	23	42			
any children (%)	27	45			
Observations	295				

^a Average characteristics of the population of job seeker allowance claimants in Edinburgh over the 6 months of the study. The numbers are based on NOMIS statistics, conditional on unemployment duration up to one year.

^b Based on the baseline survey performed in the first week.

^c High educated is defined as a university degree.

Participants were asked to search for jobs using our job search engine (described later in this section) for a minimum of 30 minutes.²¹ After this period they could continue to search or use the computers for other purposes such as writing emails, updating their CV, or applying for jobs. They could stay in our facility for up to two hours. This division was useful to inform them of artificial vacancies that they had saved once they ended the search phase, i.e., before they engaged in a real application. We did not want to inform them directly when they save a vacancy, as that might alter their behavior towards the "twin" vacancy. They could also obtain a record of their saved vacancies which the Job Centres had agreed to accept to evidence part of their job search activities. So in principle their job search with us could be used as a substitute to search on the government website. 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.

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

²¹This length is unlikely to largely alter overall job search activities on which participants spent around 12 hours a week on average.

clothing voucher for job market attire for participating in 4 sessions in a row. These were discussed with the local job centres to be permissible compensation that does not constitute income.

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 which resembles the setup of most job hubs in Edinburgh that provide free access to computers to job seekers. Panels separate desks to grant privacy and 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, which includes consent to the research vacancies (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 first asked to fill in an initial survey. From week 2 onwards, they only had to complete a short weekly survey asking about job search activities and outcomes.²² After the survey they were directed to our job search platform.

3.4 Job search platform

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.

On the main job search interface participants can search using various criteria such keywords, occupations, location, salary and preferred hours, but do not have to specify all of these (see Figure 5 in the appendix for a screen shot). 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 (if specified), 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). After the latter they can indicate why they are not interested from a list of suggested answers, and either option then redirects them to the job listings where they had left off. As in most job search engines, they can modify their search criteria at any point and launch a new search.

From week 4 onward, half of the participants were offered to use an “alternative” interface which was designed to investigate how occupational breadth of 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. (2016). The key goal of the alternative interface was

²²We received no additional information about the search activities or search outcomes from the official 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.

to offer suggestions to job seekers about occupations that might be of interest to them. This was achieved by 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 selected which suggestions they found interesting after which a search was performed over all selected occupations. Of those who are offered this interface, take up is around half. Even though the alternative interface affects individual job search behavior, it is orthogonal to the randomized set up of the artificial vacancies on which this paper focuses, which is important for the validity of the empirical strategy.

4 Empirical analysis

Our empirical analysis focuses mostly 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? We will then present an analysis of the complementary survey on how the vacancies are perceived. Before doing so, we briefly describe the outcome variables we use in the analysis.

4.1 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”). At the end of the session the list of saved vacancies would be shown (which could also be accessed from home by logging in to the system). In case the list contained artificial vacancies (s)he would, at this point, be informed about the nature of these vacancies.

Our main analysis focuses on the decision to save a vacancy. This is a clear signal of interest in the job and the closest proxy of applying to the job as almost one-third of all saved (real) vacancies is eventually applied to. Of all artificial vacancies, 42% is never saved (134 vacancies), 38% is saved between 1 and 3 times (123 vacancies), and 20% is saved more than 3 times (65 vacancies). The mean number of saves is 1.9 (the full distribution is shown in Figure 6 in the appendix). As a robustness check we also present all analysis using the decision to view a vacancy as the outcome, which is already a relevant expression of interest given that the wage was announced in the listing. Of all artificial vacancies, 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 is 3.6.

4.2 Do higher wages generate more interest in vacancies?

As a first step we report how wages attract job seekers amongst *real* vacancies that were posted during our study. We observe how often each of these vacancies was saved by participants in the study. Since the number of saves is a count variable, we perform a Poisson regression on the logarithm of the offered wage. We include a subset of all vacancies that fulfills the following requirements: (1) the vacancy is posted in the Edinburgh area, (2) it has a wage announcement (3) the wage is annual.²³ Results are presented in the appendix in Table 12. For these real vacancies a higher wage is associated with significantly less saves (column (1)). The association remains (though with slightly smaller magnitude) when controlling for occupation fixed effects (column (2)). Additional controls for a temporary contract, for part time jobs, for not listing a company name and for the posting month are all highly significant, but hardly change the negative wage coefficient (columns (3)-(6)). When analyzing jobs that post an hourly wage, we find a similar negative wage coefficient that is slightly smaller in magnitude (column (6)). Finally, also a simple log-log regression leads to very similar results (column (7)).²⁴ The significantly negative relation between the wage and the number of interested job seekers is in line with the findings of Marinescu and Wolthoff (2014) who analyse jobs from Careerbuilder.com (when not controlling for job titles) as well as the results of Faberman and Menzio (2016).

Clearly, the posted wage of real vacancies is highly correlated with other characteristics. Therefore, we consider our *experimental* vacancies, in which we randomly assigned the wage. We find that the mean number of times that the vacancy with the lower wage in the pair was saved is 1.73, while the mean number of times that the vacancy with the higher wage in the pair was saved is 2.09. A two-sided paired t-test rejects that the means are equal with a p-value of 0.02. Thus we indeed find that vacancies with a higher wage attract more saves. The same conclusion can be drawn based on views: the lower wage vacancies are viewed an average of 3.23 times, while the higher wage vacancies are on average viewed 3.87 times. This difference is significant with a p-value below 0.01.

To exploit the variation in the magnitude of the wage changes we perform a regression analysis, in which the number of saves (S) is regressed on the percentage change in the wage (Δw). To exploit the pair structure of the data, we include pair fixed effects (γ_p). Since our outcome variable is a count variable, we estimate a Poisson regression model:

$$S_{ip} = \exp(\alpha + \gamma_p + \beta \Delta w_{ip} + \epsilon_{ip}) \quad (1)$$

Vacancies are indexed by subscript i and vacancy pairs by subscript p . The parameter of interest, β , can easily be transformed to measure the percentage effect of a one percent increase in salary on the number of saves, which is the wage elasticity. 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.²⁵ To

²³The first restriction is to make the analysis comparable to our experimental results. The second restriction removes 54 % of the vacancies as they report no salary. The third restriction is used to prevent a problematic comparison of hourly, daily, weekly, monthly and annual wage announcements. Of all vacancies that report a salary, 54 % reports an annual wage and thus we focus on this category.

²⁴When using viewing a vacancy as the outcome (instead of saving a vacancy) we find very similar results, see Table 13 in the appendix.

²⁵Due to the restriction on the number of artificial vacancies, we posted somewhat less vacancies during the second

Table 2: Effect of wage change on number of saves/views

	Saves			Views		
	(1)	(2)	(3)	(4)	(5)	(6)
Salary difference (in %)	0.70** (0.44)	0.69* (0.45)	0.92*** (0.43)	0.70** (0.35)	0.71** (0.36)	0.86*** (0.29)
Appearing first			0.58*** (0.13)			0.50*** (0.075)
Pair fixed effects	yes	yes	yes	yes	yes	yes
Postal code f.e.	no	yes	yes	no	yes	yes
N	240	240	240	304	304	304

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Clustered standard errors (by pair) in parentheses. All regressions are Poisson models where $\exp(\beta) - 1$ is reported (which is the percentage effect).

correct for the correlation between these four vacancies, we cluster standard errors at the pair level.²⁶ As an extension to this simple specification we add additional controls for the geographical location of the job and for the posting order.

Estimation results are presented in Table 2. We report $\exp(\beta) - 1$, which is the percentage change in saves due to a 1-percent increase in the wage (the elasticity). In column (1) we find a highly significant positive elasticity of 0.70. 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 column (2) we add 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 estimate for the salary difference coefficient. In column (3) we additionally include a dummy equal to one for vacancies that appeared first in the search results due to having the later posting date within the pair. The difference in posting date was usually one or two days, but we find that it has a significant impact on the number of views. The posting dates were assigned independent of the wage, and as expected we find that the wage coefficient only changes slightly. This indicates that our elasticity measure is relatively robust and constitutes our main finding for this section.

Our elasticity estimates are very similar in magnitude to the results of Dal Bó et al. (2013), who report that a 33 % increase in wages offered by local governments in Mexico led to a 26 % increase in show up at an assessment centre (which implies an elasticity of 0.79). They are also close to the findings of Wolthoff (2014), who report that (when controlling for job titles) a 10 % increase in wage is associated with a 7.4 % increase in applications. This seems to indicate that the elasticity is robust to the location, underlying occupation, and empirical technique which seems remarkable. Abebe et al. (2017) find a lower elasticity of assessment centre attendance with respect to the wage of around 4.5% for the clerical positions they advertise in Ethiopia, but theirs are three-month temporary jobs and a similar wage increase for permanent jobs would lead to larger reactions if job seekers trade off fixed

wave. As a result not all artificial vacancies were used twice.

²⁶The fact that each job seeker in the study might save vacancies from different pairs can create correlation between the pairs. There is no straightforward way to correct for this, but one approach is to group vacancies that are 'similar' and thus are likely to be of interest to the same job seekers and cluster standard errors at this group level. We use the two-digit occupational code (SOC) of the vacancies to do so. However, standard errors clustered at this level are actually smaller, and thus we prefer to be conservative and do not report these results.

application costs with the net present value of wages on the job.²⁷

Rather than using saving a vacancy as the outcome variable, we can also consider viewing a vacancy (which is clicking on the vacancy appearing in the listing to view detailed information). The results are presented in columns (4)-(6) in Table 2 and are very similar to the findings for the effect on saves. The estimated wage elasticities are statistically significant and around 0.7 - 0.9 depending on the exact specification.

4.3 Is the reservation wage property satisfied?

The previous section documents that a higher wage for a given vacancy induces significantly more interest from job seekers. Simple homogeneous agent models of competitive/directed search rely on this as the main reason why firms offer attractive wages. It is usually considered inconsistent with pure versions of random search. Nevertheless, slight variations of random search that are equivalent in terms of eventual hiring decisions might also be consistent: Assume that workers encounter wage offers by firms randomly as in McCall (1970) or Burdett and Mortensen (1998), but workers only bother to send a formal application if the wage is above their reservation wage. A wage range then increases the number of applications if workers differ in outside options. The main feature that makes it comparable to a random search model is that workers only reject jobs they already know they would not take. In such a model the usual reservation wage strategy that determines job acceptance also applies to job applications: If a worker encounters two identical offers and applies to the low one he will surely also apply to the high one, as that one also meets the reservation wage.

Competitive or directed search models differ in their assumption that workers can only apply to less options than they find attractive - for example because application costs are too high (e.g., Kircher (2009)). So workers have to be selective even among attractive options. In particular, these models tend to imply that workers do not send all their applications to the highest wages. Rather they sometimes go for low wages instead of a high wage because the competition for the high wage jobs would be too tough.

The second part of the empirical analysis will therefore focus on whether the reservation wage property holds. If it holds, a job seeker who sees both vacancies and shows interest in the low wage vacancy should also be interested in the high wage vacancy. We study this focusing on individuals' decisions regarding vacancies from the same pair. Again, we proxy applications by looking at saves and views.

We document two statistics relating to the relationship between saving one or both of the vacancies in a pair. We will show the probability of not saving the high-wage vacancy, conditional on saving the low wage vacancy and vice versa:

$$P(S_h = 0 | S_l = 1) \tag{2}$$

$$P(S_l = 0 | S_h = 1) \tag{3}$$

²⁷As an alternative to the Poisson model, we can estimate a log-linear model where the dependent variable is the logarithm of the number of saves (adding a constant equal to 0.1 to handle the zeros). While we prefer the Poisson specification, results from this approach are very similar and not reported for sake of brevity. We find a significantly positive wage effect that is slightly smaller in magnitude compared to the Poisson model. Furthermore we show in Table 11 in the appendix that the results are almost identical when also Glasgow pairs are included in the analysis.

where $S_l = 1$ if the low wage vacancy in a pair was saved, and similar for S_h . These probabilities are shown in column (1) of Table 3, in the first and second row. The number of observations are given in brackets. We find that out of all individuals that save the low wage vacancy in a pair, 42% does not save the high wage vacancy. The reverse probability, not saving the low wage vacancy when the high wage vacancy is saved, is 49%. Under the reservation wage property, the first number would be zero and only the second number would be non-zero. While these numbers are significantly different from each other, they are nevertheless economically quite similar and certainly both highly significantly different from zero.

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 (s)he may simply not observe the second vacancy.²⁸ The advantage of our experimental setup is that we can observe which vacancies a job seeker has ‘listed’ on their screen. So we can compute the above probabilities conditional on listing both vacancies:

$$P(S_h = 0 | S_l = 1, L_l = 1, L_h = 1) \tag{4}$$

$$P(S_l = 0 | S_h = 1, L_l = 1, L_h = 1) \tag{5}$$

These probabilities are listed in the column (2) of Table 3. We find that the conditional probabilities of not saving one of the two are only slightly lower for individuals that have listed both. Still in 39% of the cases that an individual saves the low wage vacancy, (s)he does not save the high wage vacancy. Note however from the almost identical number of observations, that almost all participants that save one of the two vacancies have listed both vacancies.

The statistical significance of these probabilities discussed earlier are derived from two tests. First, the test whether these probabilities of not saving are significantly different from zero relies on a one-sample proportion test for each of the saving probabilities. It shows that each of these is highly significantly different from zero (p-value always smaller than 0.001). Second, the test for differences between saving the high and low wage vacancies assesses whether $P(S_h = 0 | S_l = 1) = P(S_l = 0 | S_h = 1)$. Using a two-sample proportion test we find that the unconditional probabilities (column (1)) are significantly different (p-value=0.08). Also the probabilities conditional on listing both (column (2)) are significantly different (p-value<0.03).

4.3.1 Robustness

We would like to exclude as much as possible that the decision to only save one of the two vacancies is caused by other factors than the wage itself. We propose here a number of tests to exclude possible alternative explanations. First, we control for the timing of posting of the vacancy. Second, 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. Third, we exclude individuals that responded in an exit survey that they sometimes, often or always could identify

²⁸If, on the other hand, it occurs because one of the two vacancies does not fulfill the search criteria imposed by the job seeker, it is less of worry. In that case the choice as to save one of the two can be regarded as simply revealing preferences.

Table 3: Saving and viewing probabilities

	conditional on listing both			
			high wage appearing first	low wage appearing first
	(1)	(2)	(3)	(4)
Saving				
$P(S_h = 0 S_l = 1)$	0.42 (278)	0.39 (267)	0.16 (97)	0.54 (170)
$P(S_l = 0 S_h = 1)$	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
Viewing				
$P(V_h = 0 V_l = 1)$	0.38 (520)	0.36 (500)	0.21 (200)	0.50 (300)
$P(V_l = 0 V_h = 1)$	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

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 $P(S_h = 0|S_l = 1) = P(S_l = 0|S_h = 1)$ and similar for viewing (both are two-sided tests). ^b Note that these are tests for cross probabilities (for column (3) $H_0: 0.16=0.36$ and for column (4) $H_0: 0.54=0.61$) and similar for viewing (for column (3) $H_0: 0.21=0.34$ and for column (4) $H_0: 0.50=0.55$)

artificial vacancies. Fourth, we exclude vacancy pairs for which the geographical distance within the pair was more than 1 kilometer. Finally, we discuss possible learning effects among participants. Overall we find no evidence for any of these alternative explanations. We briefly discuss the results here, which are presented in the appendix.²⁹

Timing of the vacancy A first factor that could explain why job seekers do not save both vacancies 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. But even at these time intervals job seekers might be worried that the older vacancy is already taken. Cheron and Decreuse (2017) infers that more than a third of vacancies on a US search platform are already taken, and Albrecht et al. (2017) highlight that more than 50% of applications in the US goes to vacancies that are no more than two days old, and rationalize this in a competitive search framework where workers fear that an older vacancy has already been taken. There might also be other reasons that make more recent postings more attractive. Earlier we documented that also in our setting job seekers place a premium on the more recent job posting within a pair. This makes it particularly striking if job seekers would forgo both a more recent and higher-paying vacancy within the pair to choose an older and lower-paying one.

To investigate this, we compute probabilities similar to (4) and (5), but condition on which vacancy

²⁹We also check that lack of interest for some vacancies is not caused by individuals ignoring vacancies with a 40 % higher wage because they find them suspicious. In Table 14 in the appendix we show that the probabilities are very similar when excluding all vacancy pairs with a 40 % difference.

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 3. We find that, as expected, the likelihood of saving increases when the vacancy appears first (due to being posted more recently). For those that save the low wage vacancy while the high wage one appears first, 16 % does not save the high one. This percentage is considerably lower than the 39% in column (2), however it is still a non-trivial share. The reverse probability (not saving the low one when one saves the high one and the low one appear first) is 36 % (column (4)). Both 16% and 36% are highly significantly different from zero ($p\text{-value} < 0.001$) and also from one another ($p\text{-value} < 0.01$). As a comparison, the probability of not saving the high one when the low one appears first is significantly higher (0.54) but not significantly different from the reverse probability (0.61) ($p\text{-value} = 0.18$). To assess the effect of the wage alone, ideally we would have liked to post the vacancies at the same time, and presumably we would have obtained a fraction of people that save the low wage but not the high somewhere between the extremes of 16% and 54% that we assess here. This was not feasible in our setting as this would have appeared too unrealistic, and in absence of better information we use the average estimated effect of 39% as our benchmark.

We can compute similar probabilities for viewing vacancies. These probabilities are reported on the third and fourth row of Table 3. The probability of not viewing a high wage vacancy when one has viewed the low wage vacancy is 0.38. If, in addition, the high wage vacancy appears first on the list, this probability becomes 0.21. All of these are quite close to the corresponding saving probabilities, and lead to the same conclusion: higher wages attract more interest, but a non-trivial share of job seekers is only interested in the low-wage vacancy.

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. We have 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.

In either case it could provide an explanation for not saving the second one, which is not related to the wage and is therefore different from the hypothesis we are testing. We note that at the listing stage the vacancies look rather differently (different posting date, different wage, different post code). Therefore, the decision to view a vacancies based on the listings is unlikely to be affected. Nevertheless we find very similar results. After viewing, the very similar nature of the job description might lead some to believe that they have seen the same vacancy earlier. To investigate whether these events drive our results on savings, we remove all the viewings and savings of this type from our data. This implies removing 74 vacancy-viewings (out of a total of 1145 vacancy-viewings) or 37 vacancy-saves (out of 615 saves). The new probabilities are presented in Table 15 in the appendix, which has the same structure as Table 3. We find that removing the 74 observations has very little effect on the probabilities: the non-saving probabilities decrease slightly (row 1), while the non-viewing probabilities increase slightly (row 3). We conclude that this consideration 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.

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’. 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%) and answered the question. This implies dropping 654 vacancies viewings (57 % of the observations) or 347 vacancy saves (56 % of the observations). The resulting viewing and saving probabilities are presented in Table 16 in the appendix. The probabilities of not saving or viewing the high wage vacancy conditional on doing so for the low wage one (rows 1 and 3) are slightly lower. Only the probability of not saving the high one in the case it was posted more recently is reduced quite a bit (to 0.06). However this probability is based on a very small sample of 34 vacancy observations, and most of the reduction can be attributed to the particular sample of individuals still present in the exit interview.³⁰ A 95 % confidence interval includes values between 0.02 and 0.19 and thus the number is not significantly different from our previous estimates, while it is significantly different from zero.³¹

Geographical distance within a pair We attribute differences in interest for the vacancies within a pair to the wage differences. However, within an Edinburgh vacancy pair, the location was varied randomly within Edinburgh to make sure the two vacancies would not appear to be the same. In some cases this implied only a small difference in geographical location, however in other cases the difference can be larger. 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. Since Edinburgh has substantial density around the centre, we find that for 50% of the Edinburgh pairs the distance between the two vacancies is within easy walking distance of at most 1 kilometer, while for the other 50% the distance varies from 2 to 12 kilometers.³² 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 17 in the appendix. 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.

Listing many vacancies As we show in Table 1 the number of vacancies that individuals on average list per week is large (over 500). One may be worried that because of this large number, listing a vacancy does not guarantee that the individual actually sees and considers the vacancy. To address this concern, we compute the saving/viewing probabilities for the subsample of individuals that lists less than the median number of vacancies per week (on average). This selection cuts the average number of listed vacancies per week into half (around 250 per week). Individuals that list relatively few vacancies have

³⁰ Among the selected sample that is still present in the exit interview after three months of job search, the probability of saving a low research vacancy posted earlier but not the high counterpart when it is posted later is 0.10. There are many reasons including random variation that might account for this, but it is noteworthy that - *ceteris paribus* - individuals that apply more to high wage vacancies rather than to low wage vacancies tend to have less chances of success and are more likely to end up in the exit interview.

³¹ The confidence interval is computed using the Wilson method, as recommended by Brown et al. (2001).

³² Note that these are straight-line distances and the corresponding travel distances will be somewhat larger.

more time per listed vacancy and are more likely to actually consider each of these vacancies and thus make a conscious decision about viewing/saving each of them. The results are presented in Table 18 in the appendix. We find that the probability of not saving the high wage vacancy, while saving the low wage vacancy is somewhat lower for this group (top row). When both vacancies have been listed this probability is 0.26, while it was 0.43 in the complete sample. This is expected to some extent: the smaller the set of potentially interesting vacancies, the larger the probability of saving each particular vacancy. This is also confirmed by the lower values for the reverse probabilities (rows 2 and 4). We discuss this in more detail in section 5, where we calibrate a directed search model, as it displays this obvious feature. Moreover, if we restrict the sample further, this probability does not appear to decrease further.³³ We conclude that this finding supports our hypothesis that some individuals prefer the lower wage vacancy, as even within this sample in a quarter of the cases the high wage vacancy is not saved while the low wage vacancy is saved.

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 viewing. 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 4.3.1 that only including individuals that indicated they were unable to identify artificial vacancies makes little difference for the results.

4.4 The role of perceptions about competition and working conditions

Our results suggest that some job seekers show more interest in a low-wage job than in a high-wage job, when all other characteristics within the job posting 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, an alternative explanation could be that the higher wage is 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.

³³When restricting the sample to individuals that list less than 200 vacancies per week, we find the corresponding probability is 0.31 (as compared to 0.26). Such a selection includes only 47 job seekers and thus the results become rather imprecise.

To investigate whether this is the case, we designed and conducted an online survey. Participants were recruited through the behavioural lab of the University of Edinburgh, using the non-student pool. 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. The resulting sample of participants has an average age of 31 (with a range from 19 to 62), 74% are employed (part-time or full-time), 14% are unemployed and looking for work and 6% are self employed. Of all participants, 56% indicate to be currently looking for work (even if employed). The participants included more women (68%) than men (32%). Ideally one would have collected such data on a similar sample as the one used in the field experiment and among job seekers in particular. Since recruiting job seekers is in itself an ambitious and difficult enterprise, we opted for the more convenient subject pool of the laboratory of the University of Edinburgh. We did attempt to exclude students from the survey, as they might have less experience on the labor market (only 2.5 % of our sample reported to be a student).

Each survey participant is presented with a random set of 20 of our artificial vacancies, with 3 questions about each vacancy. The questions are as follows:

1. Given the skill and experience requirements described in the job announcement (if any), how good would you expect an applicant needs to be in order to be considered for this job?
2. For someone with the skill and experience requirements described in the job announcement (if any), how much competition would you expect for this job relative to other jobs in the same profession and area?
3. For someone with the skill and experience requirements described in the job announcement (if any), how would you expect the overall (non-monetary) working conditions of this job to be? Examples of non-monetary working conditions are working hours, career prospects, demands associated with the job, health and safety, etc.

Questions (1) and (2) directly relate to directed search models, which predict that higher wage vacancies should attract more interest and from better workers (where the first implication arises even in the simplest homogeneous-worker models, while the second requires heterogeneous-worker models in which in principle only one of the two channels could be active). Question (3) relates to the theory of compensating differentials.

Each question was multiple choice with five options: (1) Very much above average; (2) Above average; (3) Average; (4) Below average; (5) Very much below average. For the complete survey questions see the online appendix OA.3. 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 were 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 with substantial agreement among the responses.³⁴

³⁴To assess the quality of the measurements of perceptions about the vacancies, we compute Cronbach's alpha for

Table 4: Wilcoxon paired signed-rank test comparing

	High wage response higher	Low wage response higher	Equal response	p-value	Obs.
Q1 (Required quality)	379	211	620	<0.001	1210
Q2 (Competition)	479	230	501	<0.001	1210
Q3 (Working conditions)	374	221	615	<0.001	1210

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 a Wilcoxon signed rank test. Results are presented in Table 4. 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 require better applicants) and will thus be harder to get. These results also reject the alternative hypothesis that high wages compensate for worse working conditions (compensating differentials), which would have led to a reverse result on question 3. Instead, it seems consistent with a view in which firms compete for workers through better salaries (usually the only channel in directed search models) as well as through better working conditions. Thus, the survey results suggests that lack of interest in high-wage vacancies might be caused by higher expected competition for the job.

Next, we investigate whether the differences in perceptions about vacancies influence our positive wage elasticity estimate. To do so, we compute the average responses across participants for each vacancy, using linear values ranging 1-5 for the response choices (where a higher number means a “higher” response). Before averaging across participants, we standardize each response by subtracting the participant’s mean response to the question, and dividing by the standard deviation of the participant’s responses to the question. We do so to correct for heterogeneity across survey participants. Next, we include the mean standardized responses as control variables when regressing log-saves on the salary difference with a specification that is otherwise identical to the one used in Section 4.2. Our hypothesis is as follows. Since higher wages are perceived to indicate (1) higher required quality, (2) higher competition for the job and (3) better working conditions, we expect the inclusion of the first two to have a negative impact on the number of saves/views *and* controlling for these should lead to a higher salary coefficient while the third perception (working conditions) should have a positive coefficient and lead to a reduced salary coefficient.

The regression results are presented in Table 5. Column (1) reproduces the main finding from Table 2. In column (2) we add the perceptions and find that only the perception about working conditions is significant and positive. The quality (question 1) and competition (question 2) perceptions have,

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 9. 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 9), we find that agreement is particularly high for question 1 and 3, while somewhat lower for question 2.

Table 5: Effect of salary on saves/views, controlling for perceptions

	Saves		Views	
	(1)	(2)	(3)	(4)
Salary difference	0.92*** (0.43)	0.74* (0.58)	0.86*** (0.29)	0.74** (0.37)
Q1 (quality) standardized		-0.053 (0.18)		0.13 (0.17)
Q2 (competition) standardized		-0.087 (0.15)		-0.16 (0.098)
Q3 (working conditions) standardized		0.35* (0.22)		0.30** (0.14)
Appearing first	0.58*** (0.13)	0.59*** (0.13)	0.50*** (0.075)	0.50*** (0.075)
Pair fe	yes	yes	yes	yes
Postal code fe	yes	yes	yes	yes
N	240	240	304	304

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Clustered standard errors (by pair) in parentheses. All columns are Poisson model regressions where $\exp(\beta) - 1$ is reported (which is the percentage effect).

as expected, a negative coefficient, but are not statistically significant. Given the significant positive coefficient for working conditions, the lower wage elasticity is in line with the hypothesis. In Table 19 in the appendix we show that including the perceptions questions one-by-one in the regression has the expected effect on the wage coefficient (except for the quality perception). In column (3) and (4) of Table 5 we find that results are very similar when considering views instead of saves.

An alternative is to look at the conditional probabilities that we analyzed in section 4.3. Thus, for those individuals that save the low-wage vacancy, we regress a dummy for also saving the high wage vacancy on the wage difference in the pair and add the perception survey results. This a more direct approach to testing whether perceptions play a role in the decision to save only the low-wage vacancy in a pair. Note that in this setup we have multiple observations for (some) individuals, and can include individual fixed effects. On the other hand, the outcome is saving the high-wage vacancy conditional on saving the low-wage vacancy and we can therefore not include pair fixed effects. To net out vacancy fixed effects, we include differences in wages and perceptions within the pair as regressors. The exact regression specification is:

$$\mathbb{1}_{(S_h)_{pj}} = \alpha + \beta \Delta w_p + \gamma_1 \Delta P_{1p} + \gamma_2 \Delta P_{2p} + \gamma_3 \Delta P_{3p} + \delta PO_p + \pi_j + \varepsilon_{pj} \quad (6)$$

Where subscript p denotes vacancy pairs and j indexes participants. $\mathbb{1}_{(S_h)_{pj}} = 1$ if individual j saved the high wage vacancy from vacancy pair p . The percentage difference in wage within pair p is denoted by Δw_p , the perception differences (for each of the three questions) between the high and low wage vacancy in pair p are given by ΔP_p . We also control for the posting order: a dummy for the high wage

Table 6: Individual level regression with $P(S_h = 1|S_l = 1)$ as outcome

	Saves		Views	
	(1)	(2)	(3)	(4)
Salary dif in the pair	-0.018 (0.31)	0.083 (0.42)	0.42* (0.24)	0.49 (0.30)
Difference Q1 (Quality)		-0.040 (0.10)		-0.022 (0.067)
Difference Q2 (Competition)		-0.066 (0.11)		-0.10** (0.052)
Difference Q3 (Working conditions)		0.066 (0.10)		0.10 (0.074)
High wage appears first	0.29*** (0.076)	0.29*** (0.076)	0.24*** (0.051)	0.25*** (0.051)
Constant	0.45*** (0.090)	0.44*** (0.10)	0.38*** (0.072)	0.38*** (0.079)
P-val joint sign. Q1-Q3		.88		.2
Individual fe	yes	yes	yes	yes
N	278	278	528	528

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (clustered by individual) in parentheses. Each observation is one individual saving/viewing one lower-wage vacancy from a pair. The outcome variable is a dummy for also saving/viewing the higher-wage vacancy from the pair.

vacancy appearing first (PO_p) and include individual fixed effects π_j . The model is estimated using all observations of individuals saving a low-wage vacancy from a pair.

Thus, this approach compares within individual variation across different vacancy pairs. We show the results in Table 6. For comparison, we show the result excluding perceptions in column (1). In column (2) we include the within pair difference in the three perceptions. While the precision of the estimates is again low (and none are statistically significant), the signs of all coefficients are in line with our hypothesis. A larger difference in higher required quality and a larger difference in competition for the higher-wage job make it less likely that an individual is also interested in saving that job, while relatively better working conditions make it more likely that (s)he is interested in the high wage vacancy. Given the lack of power, these results do not provide conclusive evidence, but they are suggestive of the idea that perceptions about the probability of getting the job play a role in the choices regarding where to apply. Note that when using views instead of saves, we have slightly more power. In columns (3) and (4) we find very similar results using views instead of saves, and competition is detected as a significant factor.

4.5 Additional explorative analysis

We restricted ourselves to no more than 2% research vacancies in the sample of real vacancies to keep the environment realistic. Since we also have a relatively limited sample of job seekers, we lack power to delve into more detailed subdivisions amongst different vacancies or carefully investigate individual heterogeneity in behavior. With this caveat in mind, we present suggestive results related to two questions: 1) Is the elasticity similar for wage increases and wage decreases?, 2) Are there individual characteristics correlated with application to high and low wage vacancies? We briefly discuss the results here, which are presented in the appendix.

4.5.1 Wage increases and wage decreases

Table 10 in appendix A.6 reports elasticities of viewing and saving distinguishing between wage increases and wage decreases. We find that the elasticity of the number of saves is significant for the wage increases where we have the majority of observations (129 pairs), and shows no statistical significance for our limited observations with wage decreases (32 pairs). The elasticity of the number of views is significantly positive for both increases and decreases.

Reductions in wages raise the problem that they might lie below the minimum wage. So any wage decreases that we did implement naturally had to be assigned to higher-wage vacancies (where the point wage or the lower bound of the wage range lies sufficiently above the minimum wage), while wage increases were also at low wage vacancies where decreases were not feasible (which constitutes 40% of our sample). When we separately evaluate the effect for low wage vacancies with wage increases, high wage vacancies with wage increases, and high wage vacancies with wage decreases (column 5) the effects are concentrated in the first set.³⁵ One reason for this finding might be that vacancies with a higher baseline salary are more likely to report a wage range rather than a single wage point. Wage ranges might be less informative to potential applicants, and thus changing these wages affects them less. We can estimate different elasticities for vacancies that report a wage range and vacancies that report a wage point, and additionally for low-wage vacancies and high-wage vacancies. We find that for low-wage vacancies the elasticity is high regardless of whether the wage is a point (1.07) or range (1.27***). For high-wage vacancies, the elasticity is smaller in general, but it seems relatively larger for point-wages (0.58) than for range-wages (0.03). Due to lack of power when estimating four distinct elasticities, most of the coefficients are not statistically significant, and neither are the differences between them.

4.5.2 Individual heterogeneity

Our results raise the question whether different types of job seekers have different responses to wage differentials. For example, whether “better” applicants are more likely to apply to higher wage vacancies. Or, whether heterogeneity across job seekers can explain wage differentials (if, for example, minorities or women are less likely to apply to higher wage vacancies). While these questions are

³⁵Rather than interacting the salary difference with an indicator for a low or high baseline salary, we can interact the salary difference with the (continuous) baseline salary. The results are similar, with a significant positive coefficient for the salary difference (1.75**) and a significantly negative interaction between the salary difference and the baseline hourly salary (-0.16*).

relevant, and we have collected several individual characteristics, our study has not been designed to answer these questions. Our sample size is too limited to identify such heterogeneity.

Keeping this in mind, we have performed several analyses concerning the relation between the response to the posted wage and individual characteristics. First, we consider the characteristics of individuals that save the lower or higher wage vacancies. We computed the *average* characteristics of all applicants that save the increased wage vacancy in a pair and the *average* characteristics of all applicants that save the unchanged wage vacancy in a pair and aggregated this across all pairs.³⁶ We test whether these two groups differ and find no significant difference in terms of gender, ethnicity (share of whites), age, education (share with university degree), family type (share with children) and unemployment duration (share with less than 6 months unemployment duration at the start of the study). Second, we take the probability of saving a high wage vacancy, conditional on saving the low wage vacancy and regress this on individual characteristics. We find that the same set of individual characteristics does not predict this outcome significantly, with the exception of age (significantly negative at the 10% level). The same holds when using views instead of saves. If, as an alternative outcome, we use the probability of saving the high wage vacancy, conditional on saving one of the two vacancies in the pair, we find that none of the individual characteristics are significant.

Thus we find no evidence of heterogeneous responses, but we emphasize that due to the small sample size we are hesitant to draw any conclusions from these findings. The imprecision of the estimates does not rule out the possibility of sizeable effects.

5 Directed search model

In the empirical section we documented two facts. First, a job that offers a high wage attracts more interest than the same job with a lower wage. Second, a non-trivial fraction of applicants saves the low wage job, but not 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. While our empirical results relate to jobs saved, we write the model in terms of applications as we are not aware of a model that makes the distinction between both.

Our attempt is to use the simplest model that still allows us to meaningfully talk about both of the above facts. An obvious candidate is the canonical directed search model with homogeneous workers and firms, workers can apply to one job per period, and in equilibrium all firms offer the same wage (as in Peters (1991) and in a special case of Albrecht et al. (2006)). Then one can assess the wage elasticity by studying how applications change when a job changes the wage offer away from the equilibrium wage. We aim to stay close to this benchmark, but need to include some additional modelling elements to do justice to our setting. We discuss these in turn to highlight exactly the role of each of them, and proceed to the formalities thereafter.

Two additional elements are needed in order to assess the second fact: most importantly we need a model with multiple job applications, otherwise it follows trivially that a worker who applies to the

³⁶We did so separately for all pairs with wage increases and all pairs with wage decreases.

low wage will not apply to the high wage since (s)he has no further application left, which would make it impossible to match the fact that some workers do apply to both wages. Models with multiple job applications include Albrecht et al. (2006), Galenianos and Kircher (2009), Kircher (2009) and Wolthoff (2014). Moreover, even though for tractability all such models use a continuum of agents, for our purpose it is useful if one can easily refine the model so that individuals only see a finite number of options. This ensures that a worker who sees two options has indeed a positive probability of applying to both, which is important to match the data but is not true in standard continuum models. Luckily, it is known since Acemoglu and Shimer (1999) that the original equilibrium when workers see all offers remains an equilibrium when all see at least one more offer than they have job applications, as long as the original equilibrium has a unique wage on the equilibrium path.³⁷ Among models with multiple job applications only the one 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.

A second modelling element is necessary to address the first fact: we need to introduce 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 worker, which places the structure of Postel-Vinay and Robin (2002) 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. We also include a chance draw that a given applicant might not be suitable for the job at hand, otherwise the observed number of job applications and vacancies imply too high a matching rate.

5.1 The model

Assume time is discrete, β is the discount factor and workers have $N \geq 2$ 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. Assume that firms also observe whether applicants are already employed elsewhere. In this case they always prefer to give the job to an unemployed worker because of the assumptions that firms compete for a worker, and competing with another employer leaves no surplus to the firm.³⁸ From the perspective of an unemployed worker, this means that unemployed workers compete only with unemployed workers for any given job.

It is well-known in the literature that a suitable worker who applies to a job with wage w where in expectation also $A\lambda(w)$ other suitable workers queue for a job has a chance of $(1 - e^{-A\lambda(w)})/A\lambda(w)$

³⁷This was proven for a single job application but the logic extends directly to multiple job applications; see below.

³⁸Employed workers still find a job more easily if working sufficiently raises the chance of being suitable for a new job.

to get a job offer (see, e.g., Wright et al. (2017) for a derivation). When sending the application the worker is not sure whether he will be suitable, so his chances ex ante are

$$m(\lambda(w)) = \frac{A(1 - e^{-A\lambda(w)})}{A\lambda(w)}. \quad (7)$$

If the unemployed worker only gets this one job offer, (s)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 (s)he gets more than one offer (s)he is paid the 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, (s)he gets unemployment benefits b . Once employed, there is a chance $1 - \delta$ of loosing the job between one period and the next. Otherwise, (s)he faces probability \tilde{m}^* of getting at least one other job offer in which case the competing firms both bid up the wage to the marginal product, and we denote by ρ the probability that (s)he ends up taking the new offer. We assume that a worker who applies today only starts on the job next period (and the separation shock only starts once (s)he is actually on the job).

For the interested reader, the full equilibrium derivation is developed in appendix A.3. Here we just represent the key equations for a worker to explain the economic forces at work. We focus on an equilibrium with a unique wage w^* and an associated job finding probability $m^* = m(\lambda(w^*))$. We also allow for a single firm that deviates and posts wage w^d which may not coincide with the equilibrium wage. Consider a worker who applies with one application to the deviant and with his other $N - 1$ applications to firms with the equilibrium wage. His utility $U(w^d)$ is given by

$$\begin{aligned} U(w^d) = & b + m(\lambda(w^d))[1 - m^*]^{N-1}\beta V(w^d) \\ & (N - 1)m^*[1 - m(\lambda(w^d))][1 - m^*]^{N-2}\beta V(w^*) \\ & + [1 - m(\lambda(w^d))][1 - m^*]^{N-1}\beta U^* \\ & + \left[\begin{array}{c} 1 - m(\lambda(w^d))[1 - m^*]^{N-1} \\ -(N - 1)m^*[1 - m(\lambda(w^d))][1 - m^*]^{N-2} \\ -[1 - m(\lambda(w^d))][1 - m^*]^{N-1} \end{array} \right] \beta V(y) \end{aligned} \quad (8)$$

where the first line reflects his unemployment benefit receipt b , as well as his probability to succeed at the deviant firm but fail at all other firms, in which case he obtains the job at wage w^d and receives continuation value $\beta V(w^d)$. The second line captures the probability that he is successful only at exactly one of the firms offering the equilibrium wage, in which case he obtains continuation value $\beta V(w^*)$. The third line captures the probability that he is unsuccessful anywhere and continues with the value βU^* from the optimal continuation search strategy in unemployment. Finally, the last line captures the probability of being successful at more than one firm, in which case the continuation wage is bid up to y and the continuation value is $\beta V(y)$.

A worker employed at wage w obtains

$$V(w) = w + \beta\delta\tilde{m}^*V(y) + \beta\delta(1 - \tilde{m}^*)V(w) + \beta(1 - \delta)U^*, \quad (9)$$

which reflects his wage payment, the probability of obtaining another job offer on the job in which case his wage rises to y , with complementary probability his wage remains constant, and finally he returns to unemployment if the job breaks up.

In equilibrium the optimal continuation value in unemployment is given by sending all applications to the equilibrium wage, so $U^* = U(w^*)$, while firms make zero profit at the equilibrium wage. The queue length at the equilibrium wage equals the ratio of aggregate vacancies to unemployed job seekers. Whenever the queue length $\lambda(w^d)$ at a deviant wages is strictly positive, it adjusts so that $U(w^d) = U^*$ since workers have to be indifferent between the equilibrium wages and the deviant wage in order for some but not all to apply there. Finally, firms do not want to deviate from equilibrium as deviant wages do not allow them to recoup the cost of posting the vacancy. For that, note that wages cannot fall below the workers' reservation wage r (devined as $V(r) = U^*$), yet for our baseline calibration this turns out not to bind at the equilibrium wage.³⁹ Once the equilibrium is solved for, one can use (8) together with the requirement that $U(w^d) = U^*$ to compute how the queue length changes when one slightly perturbs the wage away from the equilibrium wage, i.e., one can compute $\lambda'(w^*)$.

This allows us to obtain the first relevant outcome: the elasticity of worker interest with respect to the offered wage, where we here take interest to be the number of applications (which is appropriate if saving behavior in the experiment translates into applications at a constant rate). We report its value $w^* \lambda'(w^*) / \lambda(w^*)$ below, which is appropriate for local deviations. We also checked that this elasticity remains roughly constant even if we perturb the wage by 20%-40%.

Second, we are interested in the probability of applying low, but not high, conditional on observing both (for which we have an empirical estimate presented in Table 3). Calculating this probability requires setting an additional parameter X representing the number of vacancies that each job seeker observes. We derive in the appendix the probability of applying to a particular firm offering the equilibrium wage, but not to firm with a higher deviant wage, conditional on observing both. It is given by $\Pr(A_h = 0 | A_l = 1) = N(1 - p) / (N - p)$, where p denotes the probability of applying to the deviant vacancy conditional on observing it. It can be backed out from the requirement that the queue length $\lambda(w^d)$ is equal to the probability p of applying times the number of unemployed people u times the probability of seeing the vacancy captured by X/v where v is the equilibrium number of vacancies. As discussed earlier, if X equals v so that all vacancies are observed, in a large economy the probability of applying to any vacancy including the deviant becomes zero and the expression for $\Pr(A_h = 0 | A_l = 1)$ always equals one. For finite X it falls to lower levels. Finally, the probability of applying high but not low can be derived analogously.

5.2 Calibration

We calibrate the model using statistics from Edinburgh or the UK, where possible for the 4th quarter of 2013, which is the start of our experimental study. All values are listed in Table 7. We take the length of a period to be one month, and set the number of applications $N = 10.4$ based on an observed average of 2.4 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),

³⁹In Albrecht et al. (2006) this corner solution always arises at the equilibrium wage, but with on-the-job search this no longer needs to be the case.

Table 7: Calibration of parameters

Parameter	meaning	value	source
Period length		Month	
N	Applications	10.4	Based on own study
β	Discount factor	0.9914	10% discounting per year
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	0.207	NOMIS statistics on off-flow of JSA claimants (UK, 2013, Q4) imply a weekly outflow rate of between 5.0 and 6.4%. We pick 5.2% weekly, which is 20.7% monthly.
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
Replacement rate	Ratio of benefits to average wage (in the population)	0.6	UK replacement ranges from 10% to 78% (depending on previous wage and family characteristics)
\tilde{M}	Job-to-job transition rate	0.0082	ONS statistics: 722,000 job-to-job transitions quarterly (UK, 2013, Q4) while employment was 30,288,000
ρ	Probability of switching jobs conditional on an offer	1	Set to 1.
X	Number of observed vacancies	42.9	Experimental observation on number of saved vacancies

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 weekly job finding rate in the UK in the fourth quarter of 2013 between 5% and 6.4%. We pick an intermediate value of 5.2% per week, which is 20.7% per month. 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 (7).

For the separation rate we use statistics on the UK labour market from Gomes (2012) and take the sum of 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 60% 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 $w = w^*$ searches on the job, because there are no gains to moving jobs for those that already work at $w = y$. A worker receiving a competing job offer has his wage bid up to productivity, and for simplicity we assume that the worker moves to the new job. The job-to-job transition rate observed in the data (\tilde{J}) then equals $n_w \tilde{j} / (n_w + n_y)$, where n_w is the share of workers earning the low wage, n_y the share of workers earning the high wage and \tilde{j} is the matching rate for workers employed at wage w . 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 transition rate for employed workers \tilde{j} (see appendix A.4 for the derivations).

Finally, for the second question we need to take a stance on the number of observed vacancies X , and we use that the average number of vacancies the participants in our study saved was 9.9 per week. Thus we set $X = 9.9 * 4\frac{1}{3} = 42.9$. We assess the probabilities p_c (q_c) of applying to a low (high) wage vacancy but not to the high (low) wage vacancy using a salary difference between the equilibrium wage and the deviant wage of 30%.

5.3 Predictions

Now all elements necessary to compute r , w^* and $\lambda'(w^*)$ are known. We present the calibration results in the upper panel of Table 8. The queue length at jobs offering the equilibrium wage is 49. The equilibrium wage is 0.151 and slightly larger than the reservation wage (0.138).⁴⁰

The elasticity of the queue length with respect to the offered wage ($\lambda'(w^*) \frac{w^*}{\lambda(w^*)}$) is 0.96, which is very close to the corresponding empirical value that we estimated in section 4 (0.92, presented in the third column of Table 2).⁴¹

⁴⁰In this calibration, the equilibrium wage is larger than the reservation wage. This is not a result that holds in general. A slightly different calibration leads to an equilibrium wage equal to the reservation wage. Specifically, the empirical weekly job finding lies between 5% and 6.4% (NOMIS statistics) and if we pick a value of 5.5% (instead of 5.2%), we find an lower bound solution for the equilibrium wage. In that case, we find $w^* = r = 0.101$. The predictions are quite similar with this alternative calibration: a slightly lower wage elasticity of 0.52, $P(A_h = 0 | A_l = 1) = 0.74$ and $P(A_l = 0 | A_h = 1) = 0.78$.

⁴¹Note that alternatively we could compute the elasticity by picking a specific wage increase and solving equation (8) for the queue length that makes job seekers indifferent between applying to equilibrium wage vacancies and the deviant wage vacancy. This gives a very similar result for the elasticity and we find only very minor differences depending on the size of the wage increase used.

Table 8: Model predictions

	Model parameter	Model value	Experimental estimate
Equilibrium outcomes:			
Queue length	$\lambda(w^*)$	43	
Equilibrium wage	w^*	0.151	
Reservation wage	r	0.138	
Unemployment benefits	b	0.53	
Employment at low wage	η_w	6,203,086	
Employment at high wage	η_y	38,758,222	
Key predictions:			
Wage elasticity	$\frac{w^* \lambda'(w^*)}{\lambda(w^*)}$	0.96	0.92
$P(A_h = 0 A_l = 1)$	p_c	0.71	0.39
$P(A_l = 0 A_h = 1)$	q_c	0.78	0.48

Next consider the probability p_c (q_c) of applying to a low (high) wage vacancy but not to the high (low) wage vacancy. We find that $p_c = 0.71$ and $q_c = 0.78$, (see Table 8). The corresponding probabilities obtained from the experiment are 0.39 and 0.48 (using saved vacancies as a proxy for applications, see Table 3). The model overstates the overall magnitude but only slightly understates the percentage point difference. It nicely captures the fact that there does not seem to be a large difference between applying-low-but-not-high and the opposite. The predicted probabilities are highly dependent on the value of X , the number of vacancies that a job seeker observes. In order to get probabilities that match our empirical estimates closely, we would need to set $X = 20$ (which would give $p_c = 0.31$ and $q_c = 0.51$). Even though we observe in our experiment that job seekers save around 40 vacancies per month, it is not obvious that all these are really of interest. Thus, one might take the discrepancy between model and empirical findings as an indication that the actual number of ‘relevant’ jobs is smaller than the number of saved vacancies. In appendix A.5 we discuss in more detail how the choice of X affects predictions.

These results indicate that a simple parametrized directed search model generates a wage elasticity that is very close to what we observed in the experimental data. It also shows that such a model can easily replicate the fact that many individuals apply only to the low wage job but not to a high wage job even if they observe both, and that the probability of applying low but not high is in fact not that different from the probability of applying high but not low. A slightly different choice for the number of observed vacancies would have allowed us to rationalize this nearly perfectly, even though at first sight our empirical finding might have been puzzling. While we kept the setup deliberately simple, this gives hope that this model class is indeed able to capture some of the pertinent margins that drive job search behavior and are relevant for firms’ wage setting, not only qualitatively but also quantitatively.

6 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 is consistent with a view in which job seekers send only a limited amount of job applications strategically, and a calibrated directed search model with these features can generate reasonable magnitudes. The findings appear inconsistent with a view in which firms encounter applicants at random, or in which workers see a random set of job announcements and apply to all above their reservation wage.

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, real vacancies contained over 80 % of all job openings in the UK, and research vacancies were distributed across a large number of occupations but were few in number. This ensures that the behavior observed in the study is likely to be very close to real life behavior. One caveat is that our results are based on whether vacancies were viewed and saved, rather than whether actual applications were submitted. This limits effort and time spent by the participants on non-real vacancies, but it may not be a perfect predictor of actual applications. In this respect our study is similar to audit studies for resumes, where the outcome is the callback rate rather than actual job offers.

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 market. 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, especially the further study which type of workers shy away from applying to high wage jobs despite otherwise equal job descriptions. 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

- Abebe, G., Caria, S., and Ortiz-Ospina, E. (2017). The selection of talent: Experimental and structural evidence from ethiopia. mimeo.
- Acemoglu, D. and Shimer, R. (1999). Holdups and efficiency with search frictions. *International Economic Review*, 40(4):827–849.
- Albrecht, J., Decreuse, B., and Vroman, S. (2017). Directed search with phantom vacancies. Mimeo.
- Albrecht, J., Gautier, P. A., and Vroman, S. (2006). Equilibrium directed search with multiple applications. *The Review of Economic Studies*, 73(4):869–891.
- Ayres, I. and Siegelman, P. (1995). Race and gender discrimination in bargaining for a new car. *The American Economic Review*, 85(3):304–321.
- Banfi, S. and Villena-Roldán, B. (2016). Do high-wage jobs attract more applicants? Directed search evidence from the online labor market.
- Belot, M., Kircher, P., and Muller, P. (2016). Providing advice to job seekers at low cost: An experimental study on online advice. IZA Discussion Papers 10068.
- 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.
- Brown, L. D., Cai, T. T., and DasGupta, A. (2001). Interval estimation for a binomial proportion. *Statist. Sci.*, 16(2):101–133.
- Burdett, K. and Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 39(2):257–273.
- 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.
- Cheron, A. and Decreuse, B. (2017). Matching with phantoms. *Review of Economic Studies*, 84:1041?1070.
- 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.
- Faberman, R. J. and Menzio, G. (2016). Evidence on the relationship between recruiting and starting wage. Working Paper 22915, National Bureau of Economic Research.

- Fehr, E. and Goette, L. (2007). Do workers work more if wages are high? evidence from a randomized field experiment. *The American Economic Review*, 97(1):298–317.
- Galenianos, M. and Kircher, P. (2009). Directed search with multiple job applications. *Journal of Economic Theory*, 144(2):445–471.
- Goldberg, J. (2016). Kwacha gonna do? Experimental evidence about labor supply in rural Malawi. *American Economic Journal: Applied Economics*, 8(1):129–49.
- Gomes, P. (2012). Labour market flows: Facts from the united kingdom. *Labour Economics*, 19(2):165 – 175.
- Gourio, F. and Rudanko, L. (2014). Customer capital. *Review of Economic Studies*, 81:1102–1136.
- Holzer, H. J., Katz, L. F., and Krueger, A. B. (1991). Job queues and wages. *The Quarterly Journal of Economics*, 106(3):739–768.
- Hornstein, A., Krusell, P., and Violante, G. L. (2011). Frictional wage dispersion in search models: A quantitative assessment. *The American Economic Review*, 101(7):2873–2898.
- Hosios, A. (1990). On the efficiency of matching and related models of search and unemployment. *Review of Economic Studies*, 57:279–298.
- 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.
- Kroft, K., Lange, F., and Notowidigdo, M. J. (2013). Duration dependence and labor market conditions: Evidence from a field experiment*. *The Quarterly Journal of Economics*, 128(3):1123.
- Lahey, J. N. (2008). Age, women, and hiring An experimental study. *Journal of Human Resources*, 43(1):30–56.
- Leibbrandt, A. and List, J. A. (2015). Do women avoid salary negotiations? evidence from a large-scale natural field experiment. *Management Science*, 61(9):2016–2024.
- Lentz, R. and Moen, E. R. (2017). Competitive or random search. 2017 Meeting Papers 1128, Society for Economic Dynamics.
- Lester, B., Visschers, L., and Wolthoff, R. (2015). Meeting technologies and optimal trading mechanisms in competitive search markets. *Journal of Economic Theory*, 155:1–15.
- 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.
- Mas, A. and Pallais, A. (2016). Valuing alternative work arrangements. Working Paper 22708, NBER.
- McCall, J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 84:113–126.
- Menzio, G. and Shi, S. (2011). Efficient search on the job and the business cycle. *Journal of Political Economy*, 119:468–510.
- Moen, E. R. (1997). Competitive search equilibrium. *Journal of Political Economy*, 105(2):pp. 385–411.
- Mortensen, D. and Pissarides, C. (2003). Taxes, subsidies, and equilibrium market outcomes. In Edmund, P., editor, *Designing Inclusion: Tools to Raise Low-end Pay and Employment in Private Enterprise*, volume 4, Part B, pages 44–73. Cambridge University Press.
- Mortensen, D. and Wright, R. (2002). Competitive pricing and efficiency in search equilibrium. *International Economic Review*, 43:1–20.
- 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.
- Peters, M. (1991). Ex ante price offers in matching games: Non-steady states. *Econometrica*, 59:1425–1454.
- Peters, M. (1997). On the equivalence of walrasian and non-walrasian equilibria in contract markets: The case of complete contracts. *Review of Economic Studies*, 64:241–264.
- Petit, P. (2007). The effects of age and family constraints on gender hiring discrimination: A field experiment in the French financial sector. *Labour Economics*, 14(3):371 – 391.
- Postel-Vinay, F. and Robin, J.-M. (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica*, 70(6):2295–2350.
- Rooth, D.-O. (2009). Obesity, attractiveness, and differential treatment in hiring A field experiment. *Journal of Human Resources*, 44(3):710–735.
- Shi, S. (2001). Frictional assignment, part i: Efficiency. *Journal of Economic Theory*, 98:232–260.
- The New York Times (2009). 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.
- Wright, R., Kircher, P., Julien, B., and Guerrieri, V. (2017). Directed search: A guided tour. NBER Working Paper 23884.

A Appendix

A.1 Institutional Setting

In Figure 4 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 4: Aggregate labor market statistics



A.2 Job Search Interface

Figure 5 provides a screen shot of the main job search interface.

A.3 Equilibrium Derivation

In this section we derive the model equilibrium and the reaction to deviations. In equilibrium 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 . Similar to Albrecht et al. (2006) we will show that in equilibrium there is only one wage w^* . If all other firms offer this wage, an individual firm maximizes

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

Figure 5: Standard search interface

The screenshot shows a web browser window with the URL <https://www.jobsearchstudy.ed.ac.uk/index.php?r=vacancy/newSearch>. The page header includes the University of Edinburgh logo and the text 'THE UNIVERSITY of EDINBURGH'. Below the header, there is a navigation bar with 'Job search study', 'My saved jobs', and 'Logout (test4)'. The main content area is titled 'Search for Jobs' and contains a search form. The form is divided into two main sections: 'General' and 'Location and Salary'. The 'General' section has a 'Keywords' input field, two 'Occupations' dropdown menus, and an 'Hours' dropdown. The 'Location and Salary' section has 'Location' (city/postcode and radius), 'Salary' (min/max and frequency), and a checkbox for 'Include jobs with no salary information'. At the bottom of the form, there are filters for 'Jobs posted' (In the last month) and 'Order by' (Date (newest first)), along with 'Search' and 'Clear search' buttons.

where $\Pi(w)$ is the net present value of having the worker until he separates and $n(\lambda(w)) = (1 - e^{-A\lambda(w)})$ is the well-known formula for the matching rate of firms. Note that production starts next period, therefore the discounting. The firm only gets 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:

$$\begin{aligned} \Pi(w) &= y - w + \beta\delta(1 - \tilde{m}^*)\Pi(w) \\ \Leftrightarrow \Pi(w) &= \frac{y - w}{1 - \beta\delta(1 - \tilde{m}^*)}, \end{aligned} \quad (10)$$

where \tilde{m}^* is job finding probability for employed workers which triggers a bidding war between the firms that dissipates all further profit from this worker.⁴² 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

⁴²This can directly be measured from the data, but can be built assuming that 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. For workers employed at a wage below productivity the job transition rate is given by $\tilde{j}^* = 1 - \left(1 - e^{-A\lambda(w)} \tilde{A} \frac{1 - e^{-\tilde{A}\mu(w^*)}}{\tilde{A}\mu(w^*)}\right)^{\tilde{N}}$. This expression accounts for the fact that an application is only successful if no unemployed worker applies (probability $e^{-A\lambda(w)}$). 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.

hope to get two offers, in which case they can bid up the wage). However, with on-the-job search, this might not necessarily be the case and so we study the general case. In equilibrium the market utility of workers is determined by sending all N applications to firms offering w^* . 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 cases, captured by the last line, the worker gets to work at her marginal product y . This reduces to

$$\begin{aligned} U^* &= \frac{b + Nm^*[1 - m^*]^{N-1}\beta V(w^*)}{1 - [1 - m^*]^N \beta} \\ &\quad + \frac{[1 - [1 - m^*]^{N-1}[1 + (N - 1)m^*]]\beta V(y)}{1 - [1 - m^*]^N \beta}. \end{aligned} \quad (11)$$

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

$$V(y) = \frac{y + \beta(1 - \delta)U^*}{1 - \delta\beta}. \quad (12)$$

If the worker currently works at a job with wage $w = w^*$ then (s)he continues searching.

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

$$\begin{aligned} V(w) &= w + \beta\delta\tilde{m}^*V(y) + \beta\delta(1 - \tilde{m}^*)V(w) + \beta(1 - \delta)U^* \\ \Leftrightarrow V(w) &= \frac{w + \beta\delta\tilde{m}^*V(y) + \beta(1 - \delta)U^*}{1 - \delta\beta(1 - \tilde{m}^*)}. \end{aligned} \quad (13)$$

Jointly, equations (11), (12) and (13) can be solved to obtain an expression for the value of unemployment U^* as a function of w^* , m^* , \tilde{m}^* and the parameters β, δ, y, b .

The equilibrium wage is set by firms to maximize profits (equation (10)). The first-order condition

is

$$\begin{aligned} n'(\lambda(w^*))\lambda'(w^*)[1 - m(\lambda(w^*))]^{N-1}\beta\frac{y - w^*}{1 - \beta\delta(1 - \tilde{m}^*)} \\ - n(\lambda(w^*))\lambda'(w^*)[1 - m(\lambda(w^*))]^{N-1}\beta\frac{1}{1 - \beta\delta(1 - \tilde{m}^*)} \leq 0 \end{aligned} \quad (14)$$

The wage cannot fall below the reservation wage r (which is defined implicitly by $V(r) = U^*$). Thus, the first-order condition holds with equality if $w^* > r$.

Now we proceed to deviations, in order to derive the responsiveness of the number of applicants to the offered wage. Note that again only the behavior of unemployed workers is important. If a single firm deviates by offering a wage $w^d > w^*$, the unemployed applicant's return, when sending one of the applications to the deviant vacancy, is given by (8) in the main text. 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 $U(w^d) = U^*$. Applying the implicit function theorem implies that the derivative of $\lambda(w)$ evaluated at w^* is

$$\lambda'(w^*) = - \frac{m^*[1 - m^*]^{N-1}\beta\frac{1}{1 - \beta\delta(1 - \tilde{m}^*)}}{\begin{pmatrix} m'(\lambda(w^*))\lambda'(w^*)[1 - m^*]^{N-2}(1 - Nm^*)\beta V(w^*) \\ -m'(\lambda(w^*))\lambda'(w^*)[1 - m^*]^{N-1}\beta U^*(w^*) \\ +m'(\lambda(w^*))\lambda'(w^*)(N-1)m^*[1 - m^*]^{N-2}\beta V(y) \end{pmatrix}}$$

The first relevant outcome is the elasticity of worker interest with respect to the offered wage, where we here take interest as number of applications (which is appropriate if saving behavior in the experiment translates into applications at a constant rate). It is given by :

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

The second relevant outcome is the probability of not applying to the high wage vacancy conditional on applying to the low wage vacancy: $P(A_h = 0 | A_l = 1)$. Calculating this probability requires setting an additional parameter X , equal to the number of vacancies that each job seeker 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.

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^d and let p be the probability that an individual applies to this particular vacancy, conditional on observing it. The queue length at the deviant vacancy is given by p multiplied by the number of people that observe it:

$$\begin{aligned} \lambda(w^d) &= p\frac{X}{v}u \\ p &= \frac{\lambda(w^d)v}{Xu} \end{aligned} \quad (15)$$

The value p should be such that workers are indifferent between applying to the deviant vacancy and equilibrium wage vacancies. Using equation 8 we compute $\lambda(w^d)$ 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:

$$q = \frac{N - p}{X - 1} \quad (16)$$

We are interested in the probability of applying low, but not high, conditional on observing both. In this framework, conditioning on applying low means fixing one of the applications to be sent to an equilibrium wage vacancy. As a result, there are $N - 1$ applications left and the probability of *not applying* to the deviant vacancy becomes $p_c \equiv \Pr(A_h = 0 | A_l = 1) = \frac{N(1-p)}{N-p}$.⁴³ The reverse probability, not applying low conditional on applying high, can be defined as following. Assume an individual sends one application to the deviant wage, then there are $N - 1$ applications left for $X - 1$ equilibrium wage vacancies. Thus, the probability of applying to each of these is simply $\frac{N-1}{X-1}$, and the probability of not applying to one of these is $q_c = 1 - \frac{N-1}{X-1} = \frac{X-N}{X-1}$.

A.4 Derivation job-to-job transitions

Define the share of people employed at the low wage (w) as n_w . Inflow I and outflow O are:

$$\begin{aligned} I_w &= uNm^*(1 - m^*)^{N-1} = up_1 \\ O_w &= n_w(1 - \delta + \tilde{j}\delta) \end{aligned}$$

Equating inflow and outflow gives

$$n_w = \frac{up_1}{(1 - \delta + \tilde{j}\delta)} \quad (18)$$

Define the share of people employed at productivity (y) as n_y . Inflow I and outflow O are:

$$\begin{aligned} I_y &= u(1 - Nm^*(1 - m^*)^{N-1} - (1 - m^*)^N) + n_w\delta\tilde{j} = up_2 + n_w\delta\tilde{j} \\ O_y &= n_y(1 - \delta) \end{aligned}$$

Again, equating inflow and outflow gives:

$$n_y = \frac{up_2 + n_w\delta\tilde{j}}{1 - \delta} \quad (19)$$

⁴³The probability of not applying high conditional on applying to one particular low wage vacancy equals:

$$\Pr(A_h = 0 | A_l = 1) = \frac{\Pr(A_h = 0, A_l = 1)}{\Pr(A_l = 1)} = \frac{\Pr(A_l = 1 | A_h = 0) \Pr(A_h = 0)}{\Pr(A_l = 1)} \quad (17)$$

These three right-hand side terms are straightforward. We have $\Pr(A_h = 0) = 1 - p$, we have $\Pr(A_l = 1) = q$ and finally we have $\Pr(A_l = 1 | A_h = 0) = \frac{N}{X-1}$. Substituting these three into equation (17) and using equation (16) to replace q we find $\Pr(A_h = 0 | A_l = 1) = \frac{N(1-p)}{N-p}$.

Finally, the aggregate job-to-job transition rate (\tilde{J} , observed in the data) is:

$$\tilde{J} = \frac{n_w \rho \delta \tilde{m}}{n_w + n_y}$$

Substituting n_w and n_y using (18) and (19) and rewriting for \tilde{j} gives the expression for the transition rate for employed workers.

A.5 Number of observed vacancies in the model calibration

In the calibration of the directed search model (section 5.3), we have to set a value for the number of vacancies that each job seeker observes each month and from which she selects the vacancies to apply to (denoted by X). We use $X = 42.9$, based on our observation that the average weekly number of saved vacancies is 9.9. Since the model predictions depend strongly on this parameter we show here results when varying the value of X .

We observe several empirical values that could be interpreted as the set of vacancies that a job seekers observes and considers. First, the average number of weekly listed vacancies is 528. While these are all vacancies that were listed on the screen and thus observed by the job seeker, it is obvious that only a small share of these is likely to be of actual interest. Out of these, an average of 24 are viewed (clicked on to see vacancy details). This set is more likely to contain relevant vacancies, though we argue that those that were not saved are unlikely to be seriously considered. Only 9.9 vacancies are saved on average. While we can use this number (and thus set $X = 42.9$), it is not obvious that all of these vacancies are of interest, and we could assume that for example half of these are relevant vacancies, leading to $X = 21.5$. A further option is to follow the robustness check performed in Table 18, in which we only consider individuals that list relatively few vacancies and are thus more likely to consciously observe each of the two vacancies in the pair. This sample only saves 8.3 vacancies per week. And, also for this sample we could use 50% of the saves.

In Table 9 we use these various values for X and show the model predictions for the two probabilities of interest: the probability of applying low but not high ($P(A_h = 0 | A_l = 1)$) and the reverse probability ($P(A_l = 0 | A_h = 1)$). Note that most other model predictions do not depend on X (most importantly the elasticity of the number of applications) and thus we do not show those in the Table.

Table 9: Sensitivity of model predictions to choice of X (number of observed vacancies)

Weekly # of observed vacancies	Empirical source	Monthly observed vacancies (X)	$P(A_h = 0 A_l = 1)$	$P(A_l = 0 A_h = 1)$
24	# of viewed	104	0.88	0.91
9.9	# of saved	42.9	0.71	0.78
8.3	# of saved (sample that lists few)	36	0.65	0.73
4.45	50% of the # of saved	21.5	0.40	0.54
4.15	50% of the # of saved (sample that lists few)	18	0.28	0.45

The monthly values for (X) follow from multiplying the relevant weekly value with $4\frac{1}{3}$. The number of applications is $2.4 * 4\frac{1}{3} = 10.4$ in each calibration.

A.6 Additional results

Figure 6: Number of times a vacancy was saved and viewed: Edinburgh pairs




Figure 7: Number of times a vacancy was saved and viewed: Glasgow pairs




Figure 8: Cumulative distribution of the within pair difference between the higher and lower wage vacancies in the pairs




Table 10: Effect of wage change on number of saves/views: wage increases and decreases

	Saves		Views	
	(1)	(2)	(3)	(4)
Salary dif.*increases	0.96*** (0.45)		0.82*** (0.29)	
Salary dif.*decreases	0.24 (0.74)		1.84* (1.62)	
Sal. dif.*low*increases		1.21*** (0.61)		0.88*** (0.35)
Sal. dif.*high*increases		0.31 (0.46)		0.63* (0.47)
Sal. dif.*high*decreases		0.23 (0.74)		1.84* (1.62)
Appearing first	0.57*** (0.13)	0.56*** (0.13)	0.50*** (0.076)	0.50*** (0.075)
Pair fixed effects	yes	yes	yes	yes
Postal code f.e.	yes	yes	yes	yes
N	240	240	304	304

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Clustered standard errors (by pair) in parentheses. All regressions are Poisson models where $\exp(\beta) - 1$ is reported (which is the percentage effect).

Table 11: Effect of wage change on number of saves/views: Edinburgh and Glasgow pairs

	Saves			Views		
	(1)	(2)	(3)	(4)	(5)	(6)
Salary difference (in %)	0.73** (0.45)	0.72** (0.46)	0.96*** (0.44)	0.72*** (0.36)	0.73*** (0.36)	0.87*** (0.30)
Appearing first			0.58*** (0.13)			0.50*** (0.075)
Pair fixed effects	yes	yes	yes	yes	yes	yes
Postal code f.e.	no	yes	yes	no	yes	yes
N	248	248	248	318	318	318

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Clustered standard errors (by pair) in parentheses. All regressions are Poisson models where $\exp(\beta) - 1$ is reported (which is the percentage effect).

Table 12: Effect of wage on number of saved: real vacancies

	Poisson regression			Log-log regression			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log(Salary)	-0.58*** (0.010)	-0.52*** (0.046)	-0.48*** (0.046)	-0.49*** (0.045)	-0.52*** (0.049)	-0.21** (0.077)	-0.38*** (0.060)
Temporary contract			-0.21* (0.11)	-0.20 (0.11)	-0.16 (0.12)	-0.32*** (0.073)	-0.19** (0.082)
Part time			0.87*** (0.26)	0.86*** (0.26)	0.81*** (0.25)	0.21* (0.13)	0.58*** (0.098)
No company name			-0.56*** (0.050)	-0.61*** (0.044)	-0.58*** (0.048)	-0.17* (0.090)	-0.40*** (0.071)
No contacts in ad				0.68*** (0.23)	0.73*** (0.24)	0.58*** (0.14)	0.27*** (0.090)
Constant	3142.1*** (750.3)						1.71*** (0.57)
Sample	Annual wages	Annual wages	Annual wages	Annual wages	Annual wages	Hourly wages	Annual wages
Occupation f.e.	no	yes	yes	yes	yes	yes	yes
Month f.e.	no	no	no	no	yes	yes	yes
N	7173	6812	6739	6739	6739	4336	7099

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Columns (1)-(6) are Poisson model where $exp(\beta) - 1$ is reported (which is the percentage effect). Column (7) is a log-log regression where the independent variable is $\log(\text{saves}+0.1)$

Table 13: Effect of wage on number of viewed: real vacancies

	Poisson regression			Log-log regression			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log(Salary)	-0.52*** (0.0085)	-0.47*** (0.012)	-0.39*** (0.015)	-0.40*** (0.015)	-0.42*** (0.052)	-0.18*** (0.020)	-0.37*** (0.042)
Temporary contract			-0.26*** (0.021)	-0.26*** (0.021)	-0.22*** (0.085)	-0.21*** (0.022)	-0.40*** (0.046)
Part time			1.15*** (0.063)	1.14*** (0.063)	1.11*** (0.23)	0.42*** (0.036)	0.77*** (0.065)
No company name			-0.48*** (0.016)	-0.53*** (0.014)	-0.48*** (0.045)	-0.16*** (0.026)	-0.46*** (0.047)
No contacts in ad				0.55*** (0.050)	0.58*** (0.19)	0.39*** (0.036)	0.30*** (0.054)
Constant	1953.3*** (342.2)						1.92*** (0.43)
Sample	Annual wages	Annual wages	Annual wages	Annual wages	Annual wages	Annual wages	Annual wages
Occupation f.e.	no	yes	yes	yes	yes	yes	yes
Month f.e.	no	no	no	no	yes	yes	yes
N	7173	7018	6944	6944	6944	4382	7099

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Columns (1)-(6) are Poisson model where $exp(\beta) - 1$ is reported (which is the percentage effect). Column (7) is a log-log regression where the independent variable is $\log(\text{views}+0.1)$

Table 14: Saving and viewing probabilities (small difference in salary-pairs only)

	conditional on listing both			
	(1)	(2)	high wage appearing first (3)	low wage appearing first (4)
$P(A_h = 0 A_l = 1)$	0.37 (134)	0.35 (132)	0.13 (39)	0.49 (93)
$P(A_l = 0 A_h = 1)$	0.46 (170)	0.46 (163)	0.63 (90)	0.27 (73)
$P(V_h = 0 V_l = 1)$	0.38 (265)	0.37 (260)	0.25 (97)	0.48 (163)
$P(V_l = 0 V_h = 1)$	0.47 (322)	0.46 (309)	0.60 (181)	0.29 (128)

Table 15: Saving and viewing probabilities excluding participants that indicate “it is posted twice/already viewed” after having saved the first vacancy

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

Number of observations in brackets

Table 16: Saving and viewing probabilities including only participants that respond in exit survey that they could never or rarely identify artificial vacancies

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

Number of observations in brackets

Table 17: Saving and viewing 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)
$P(S_h = 0 S_l = 1)$	0.46 (139)	0.43 (135)	0.14 (41)	0.61 (94)
$P(S_l = 0 S_h = 1)$	0.52 (171)	0.51 (158)	0.65 (90)	0.35 (68)
$P(V_h = 0 V_l = 1)$	0.35 (264)	0.33 (256)	0.23 (96)	0.43 (160)
$P(V_l = 0 V_h = 1)$	0.46 (327)	0.44 (306)	0.55 (176)	0.32 (130)

Number of observations in brackets

Table 18: Viewing and saving probabilities excluding participants that listed more than the median number of vacancies per week (on average)

	conditional on listing both			
			high wage appearing first	low wage appearing first
	(1)	(2)	(3)	(4)
$P(S_h = 0 S_l = 1)$	0.30 (116)	0.26 (110)	0.11 (47)	0.40 (63)
$P(S_l = 0 S_h = 1)$	0.45 (146)	0.42 (135)	0.45 (75)	0.38 (60)
$P(V_h = 0 V_l = 1)$	0.35 (207)	0.32 (194)	0.13 (81)	0.49 (113)
$P(V_l = 0 V_h = 1)$	0.47 (253)	0.45 (235)	0.51 (134)	0.38 (101)


Table 19: Effect of salary on saves, controlling for perceptions

	(1)	(2)	(3)	(4)	(5)
Salary difference	0.92*** (0.43)	0.74* (0.58)	0.85** (0.47)	0.99** (0.65)	0.53 (0.41)
Q1 (quality) standardized		-0.053 (0.18)	0.068 (0.20)		
Q2 (competition) standardized		-0.087 (0.15)		-0.021 (0.16)	
Q3 (working conditions) standardized		0.35* (0.22)			0.29* (0.19)
Appearing first	0.58*** (0.13)	0.59*** (0.13)	0.58*** (0.13)	0.58*** (0.13)	0.58*** (0.13)
Pair fe	yes	yes	yes	yes	yes
Postal code fe	yes	yes	yes	yes	yes
N	240	240	240	240	240


* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Clustered standard errors (by pair) in parentheses. All columns are Poisson model regressions where $\exp(\beta) - 1$ is reported (which is the percentage effect).

Figure 9: Cronbach alpha's for the 16 vacancy sets


(a) All questions



(b) Question 1



(c) Question 2



(d) Question 3






Figure 10: Cronbach alpha's for the 161 vacancy pairs


(a) All questions



(b) Question 1



(c) Question 2



(d) Question 3




Table 20: Effect of salary on views, controlling for perceptions

	(1)	(2)	(3)	(4)	(5)
Salary difference	0.86*** (0.29)	0.74** (0.37)	0.63** (0.31)	1.10*** (0.45)	0.47** (0.29)
Q1 (quality) standardized		0.13 (0.17)	0.22 (0.18)		
Q2 (competition) standardized		-0.16 (0.098)		-0.076 (0.11)	
Q3 (working conditions) standardized		0.30** (0.14)			0.28** (0.13)
Appearing first	0.50*** (0.075)	0.50*** (0.075)	0.49*** (0.078)	0.51*** (0.075)	0.49*** (0.075)
Pair fe	yes	yes	yes	yes	yes
Postal code fe	yes	yes	yes	yes	yes
N	304	304	304	304	304

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Clustered standard errors (by pair) in parentheses. All columns are Poisson model regressions where $\exp(\beta) - 1$ is reported (which is the percentage effect).

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 compelled 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.

Search for Jobs

You should spend at least 30 minutes searching for jobs in the lab, after which you will be able to view / print / apply for your saved vacancies and use the rest of the computer, you have been searching for 30 minutes.

Search for jobs by entering one or more search terms below.

General	Location and Salary
Keywords <input type="text" value="Keywords (e.g. nurse)"/>	Location <input type="text" value="Enter city or postcod"/> <input type="text" value="radius"/>
Occupations <input type="text" value="Select a category"/> <input type="button" value="v"/>	Salary <input type="text" value="min"/> to <input type="text" value="max"/> <input type="text" value="Select a"/> <input type="button" value="v"/>
<input type="text" value="Select a category then an"/> <input type="button" value="v"/>	<input type="text" value="min"/> to <input type="text" value="max"/> <input type="text" value="Select a"/> <input type="button" value="v"/>
<input type="text" value="Select a category"/> <input type="button" value="v"/>	<input type="checkbox"/> Include jobs with no salary information
<input type="text" value="Select a category then an"/> <input type="button" value="v"/>	choose up to 2 salary ranges
Hours <input type="text" value="Select desired hours"/> <input type="button" value="v"/>	
choose up to 2 occupations or categories	

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.

OA.3 Vacancy perceptions survey

VACANCY SURVEY

Thank you for participating to this survey. We will show you 20 job advertisements and ask you to answer questions about these ads. After that, we will ask you a few questions about your background characteristics.

Preamble

Please confirm that you are eligible to participate to the study (click all that applies):

1. I am currently living in Edinburgh
2. I am not a student
3. I am a registered participant of the BLUE subject pool

[Participants will be shown 20 vacancies in total, one vacancy at a time, and will be asked to answer the following questions]

1. Given the skill and experience requirements described in the job announcement (if any), how good would you expect an applicant needs to be in order to be considered for this job?
 - a. Very much above average
 - b. Above average and higher
 - c. Average and higher
 - d. Below average and higher
 - e. Very much below average and higher
2. For someone with the skill and experience requirements described in the job announcement (if any), how much competition would you expect for this job relative to other jobs in the same profession and area?
 - a. Very much above average
 - b. Above average
 - c. Average
 - d. Below average
 - e. Very much below average
3. For someone with the skill and experience requirements described in the job announcement (if any), how would you expect the overall (non-monetary) working conditions of this job to be? Examples of non-monetary working conditions are working hours, career prospects, demands associated with the job, health and safety, etc.
 - a. Much better than average
 - b. Better than average
 - c. Average
 - d. Worse than average
 - e. Much worse than average
4. What is your gender?
 - a. Female
 - b. Male

5. How old are you? [text]
6. What is your current occupation?
 - a. Employed part or full time
 - b. Self employed
 - c. Unemployed and not looking for work
 - d. Unemployed and looking for work
 - e. Retired
 - f. Student
 - g. Other
7. Are you currently looking for work?
 - a. Yes
 - b. No
8. What is your unique BLUE id?
9. What is your e-mail address (that was used to send the survey link)?

OA.4 Example vacancy pair

Note that the visualizations below have been generated from the underlying data of the vacancies. Thus, this is not exactly how they appeared to the job seekers. Of course the id-numbers were not shown to the job seekers.

id: 1940173

id2:102278128

Job Title: Room Attendant

Location: Edinburgh (EH23DT)

Skills requirements: Punctual, Attention to detail, Efficient

Salary:

min: £9.66 (Per Hour)

max: £9.66

Contract: Full time

Room attendant required. Should be dependable, quality focussed and able to provide good customer service. Duties include: maintaining hotel rooms by cleaning, dusting, vacuuming, and polishing; providing linen services; other cleaning activities as required.

id: 1940172

id2:102278128

Job Title: Room Attendant

Location: Edinburgh (EH75DW)

Skills requirements: Punctual,Attention to detail,Efficient

Salary:

min: £6.90 (Per Hour)

max: £6.90

Contract: Full time

Job Purpose:

Maintains hotel rooms by cleaning, dusting, vacuuming, and polishing; providing linen services.

Duties:

* Maintains cleaning schedule priorities by following room assignment list; servicing rooms requesting early cleaning first. Skills/Qualifications:

Dependability, Quality Focus, Customer Service

id: 1870132

id2:102290576

Job Title: Welder

Location: Edinburgh (EH141AR)

Skills requirements:

Salary:

min: £11.90 (Per Hour)

max: £12.60

Contract: Full time

Temporary Assignment - Immediate start. We are looking for experienced fabricator/ welders for our client. Candidates applying for this role must have experience of MIG/TIG welding, and be able to set up and run machines in a fabrication environment. You must also have a good level of English and be able to read drawings and instructions.

Successful candidates must have an engineering background.

id: 1870133

id2:102290576

Job Title: Welder

Location: Edinburgh (EH31DT)

Skills requirements:

Salary:

min: £8.50 (Per Hour)

max: £9.00

Contract: Full time

Position available for an experienced fabricator/ welders.

Candidates applying for this role must have experience of MIG/TIG welding, and be able to set up and run machines in a fabrication environment. A good level of English and the ability to read drawings and instructions is also necessary.

Successful candidates must have an engineering background. (This is a temporary assignment, with immediate start)