How wage announcements affect job search - a field experiment

Michèle Belot, Philipp Kircher, and Paul Muller*

February 14, 2022

Abstract

In a field experiment, we study how job seekers respond to posted wages by assigning wages randomly to pairs of otherwise similar vacancies in a large number of professions. Higher wages attract significantly more interest. Still, a non-trivial number of applicants only reveal an interest in the low wage vacancy. With a complementary survey, we show that external raters perceive higher wage jobs as more competitive. These findings qualitatively support core predictions of theories of directed/competitive search, though in the simplest calibrated model, applications react too strongly to the wage. We discuss extensions such as on-the-job search that rectify this.

Keywords: Wage posting, wage competition, online job search, directed search, field experiments.

JEL-codes: J31, J63, J64, C93

**Affiliations*: Michèle Belot, Cornell University and University of Edinburgh (mb2693@cornell.edu); Philipp Kircher, Cornell University, Universite Catholique de Louvain and University of Edinburgh (pk532@cornell.edu); Paul Muller, Vrije Universiteit Amsterdam and Tinbergen Institute (p.muller@vu.nl). 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. All data used in this article is accessible from Belot et al. (2021). This study was not pre-registered as this was not common practice at the time it was carried out. The application for ethical approval related to two studies conducted on the same sample, one of them already published Belot et al. (2019)) and the present study. We describe the protocol and outcomes of interest in that ethics application. Ethical approval for the combined study was initially granted by the IRB of the London School of Economics (11/30/2011) and reaffirmed by an IRB review of the University of Edinburgh (6/4/2013). The vacancy survey described in Section 5.2 obtained separate IRB approval from the University of Edinburgh (10/23/2013).

"In conventional market situations [..] wage setting occurs prior to search so that firms can influence the allocation of resources in the market." Hosios, 1990, p. 280.

1 Introduction

How labor supply varies with the wage offer by firms has been studied by economists for a long time, and job search frictions are a key ingredient to explain and model non-competitive elasticities.[1] In particular, models of directed/competitive search, which have become increasingly popular in macroeconomics, rely crucially on the idea that higher wage postings attract more applicants, generating a wage elasticity of *job applications* (see Wright et al. (2019) for a survey of directed search). In such models workers trade off a higher wage against increased competition for the job. Given that they cannot apply to all jobs, they apply more to high wage jobs up to the point where the increased competition from other workers offsets the gain. Higher wages should attract more interest and should therefore be perceived as more competitive to get, which should induce at least some workers to refrain from applying to these high wage jobs, according to these models. This channel is crucial to the efficiency and tractability of these models (e.g., Hosios (1990), Moen (1997), Acemoglu and Shimer (1999), Mortensen and Pissarides (2003), Mortensen and Wright (2002), Shi (2002), Shimer (2005), Menzio and Shi (2011), Gurio and Rudanko (2014), Schaal (2017), Cooley et al. (2020)).

The aim of this article is threefold. First, we aim to understand whether higher wages indeed raise the number of job applications, and if so, by how much. In the environment that we study, only around half of the jobs provide wage information. One could suspect that the others omit to do so because such announcements might simply be ineffective cheap-talk that is disregarded by workers. So we concentrate on vacancies that do announce wages and investigate whether altering the wage would allow firms to affect job applications. Second, we investigate whether there are indeed indications that some workers forgo applying to higher-wage jobs, and whether the competition for such jobs might be a rationale for this. Finally, we study whether the magnitudes of the effects are in line with a simple directed search model.

In contrast to the assertion in the quote above from one of the most cited articles in the search literature, we are not aware of studies based on detailed micro-data that clearly demonstrate evidence that wage announcements are successful in affecting the allocation of labor.[2] Despite the simplicity and importance of the question, 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. This also makes it hard to study why some workers do not apply to better wages, because again the other characteristics across job advertisements might differ. Making job postings comparable remains a challenge, as discussed e.g. in Faberman and Menzio (2018). The best current control is the job title (Marinescu and Wolthoff (2020)), but notwithstanding such controls the wage predicts job search behavior, even when the wage information is withheld from potential applicants (Banfi and Villena-Roldan (2019)). This suggests that job seekers condition on other parts of the job

[1] See for example Boal and Ransom (1997), Manning (2011) and Manning (2021) for surveys.

[2] Hosios (1990) is mostly known for studying efficiency under random search, though an important section containing the quote above contains an analyses of directed search that seems largely overlooked in our profession.

description correlated with the wage. How economically relevant this is remains unanswered unless one can isolate the wage effect.

Our main empirical innovation in this regard is to twist the "audit study" methodology. This methodology has been extensively used to understand employer behavior by exposing them to fictitious resumes of job seekers which only differ in one dimension such as ethnicity or unemployment duration (e.g., Bertrand and Mullainathan (2004), Kroft et al. (2013), Farber et al. (2016)). Here we flip this idea to study the worker side of the market. We randomly assign different wages to fictitious vacancies to understand how unemployed job seekers distinguish between private sector jobs. These fictitious 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 that - all else equal - a higher wage attracts significantly more interest, with an elasticity between 0.7 and 0.9. This is in line with some estimates of labor supply elasticities on the extensive margin as we discuss below, but it particularly focuses on the search decisions of workers (rather than job acceptance and retention decisions once both sides meet).

Through the lens of directed/competitive search models, this arises because (1) workers have less applications available to them than there are jobs that interest them and therefore they simply cannot apply to all jobs above some reservation wage, and (2) jobs that pay higher wages are perceived as more competitive which is precisely the reason why workers send applications even to lower-paying jobs.[3] We provide secondary evidence on both dimensions: a non-trivial share of workers who consider applying to a low wage job do not do so for a "twin" high wage job, and external raters indeed view higher wages as more competitive. To give some magnitude, a quarter of those workers who apply to the low wage vacancy does not apply to the high wage twin, even if the high wage vacancy was posted more recently and was displayed to them. This is significant relative to the null hypothesis that workers follow the simple strategy of applying to any job with a wage above their reservation wage: i.e., among otherwise identical jobs a worker who applies to a low wage also applies to any higher wage. Finally, we confront one of the simplest directed search models with these empirical findings. It misses them by an order of magnitude. We discuss a number of potential reasons, and formally show that a simple extension to on-the-job search brings the model predictions to the right ballpark.

Akin to previous audit studies on the employer side, our empirical identification of job seeker behavior rests on randomly assigning attributes (here: wages) to fictitious entities (here: vacancies). It is practically infeasible to do this with real jobs across a large range of occupations.[4] This raises ethical issues, which are amplified in our setting because job seekers are more vulnerable than employers. In contrast to audit studies on employers which are conducted without their consent, our Ethical Review Board would not authorize such a study design without prior consent from the job seekers. For this reason, we opted for a medium scale field experiment.

We created our own job matching platform, providing access to continuously updated vacancies that were 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

[3]While there are many versions of directed/competitive search models, to our knowledge they all share these two features.

[4]As we discuss in the literature section, jobs in one specific occupation can be studied, as in Dal Bó et al. (2013) who focus on the recruitment of community developing agents in rural Mexico.

job seekers to participate in our study at the experimental laboratory in Edinburgh once every week for a duration of 12 weeks. Per session, they spent at least half an hour searching for jobs that they could save to apply later, and they could spend up to two hours to actually apply for jobs or they could 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. 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, answer clarifying questions, 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 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 in order to apply later, which for actual vacancies is a strong indicator of interest in the vacancy and likelihood of applying.[5] 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 question, that is, we did not expose that the research vacancies were intended to measure responsiveness to posted wages. In practice, these vacancies 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 main research question by assessing the difference in "saves to apply later" between the two vacancies within a pair. We also use this pair design to investigate in some more detail the application behavior towards low and high wage vacancies.

Our finding that a 1% increase in the wage results in 0.7% - 0.9% more saves accounts for pair fixed-effects; so it exploits randomization within pairs. A similar elasticity is observed for "views", i.e., whether job seekers click to obtain more information about a vacancy beyond its job title, location and wage that are displayed automatically. This is statistically significant relative to the null hypothesis that workers disregard the wage information and search randomly. This random search null hypothesis is a useful benchmark in light of the possibility that job seekers regard wage postings as non-binding in the subsequent bargaining and therefore uninformative. The rejection of this null hypothesis indicates that those parts of the labor market that use wage announcements have a significant directed search component. We nevertheless stress that our study does not rule out an important random search dimension to the overall labor market: roughly 60% of jobs in our setting do not post wages or wage ranges, and it is unclear whether these contain no wage information at all or whether other verbal indicators provide a proxy to a wage announcement as argued in Marinescu and Wolthoff (2020).

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

Further studies are needed here.[6]

We rely on the pair design to investigate another plausible hypothesis distinct from directed search: that job seekers use a simple reservation wage strategy familiar from random search and apply to all jobs that exceed it. As mentioned, in directed search models workers have to be selective as they cannot apply to all jobs that interest them. They target particular wages, trading off the wage against the difficulty of getting the job. In these models some workers will choose to apply to low wages even if an equivalent job with a higher wage is offered in the market. So they do not apply to all jobs above some reservation wage. While this is notoriously difficult to study, we provide some evidence in this vein by considering how many of those workers that save the low wage posting do not save the corresponding high wage one. This fraction is large: of those workers who save the low wage vacancy more than 40% do not save the high wage counter-part, and this still remains around a quarter even if the high wage vacancy was posted more recently.[7] This fraction remains large and significantly different from the null hypothesis of a reservation wage strategy across a large number of robustness checks, including conditioning on having been exposed to both jobs on the screen, conditioning on pairs where both jobs are located within walking distance, and among participants with short lists of jobs they are considering. While we lack power to establish which worker characteristics drive workers to only apply to low wage jobs, this remains exciting work for future studies.

We also consider the beliefs about these jobs. Our randomized experimental setup ensures that both vacancies within a pair are virtually identical in terms of all other stated dimensions of the job.[8] 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 the vacancy. Directed search theories suggest that high wage vacancies should attract more interest (and from better workers), and therefore high wage vacancies should be perceived as such. It is conceivable, though, that higher wages are viewed differently, for example as compensation for worse working conditions. 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 (A) attract more competition from other workers, (B) require an applicant to be of higher quality to be considered for the job and (C) have better non-monetary working conditions. Findings (A) and (B) support a directed search interpretation of our empirical results, while finding (C) suggests that firms compete for workers not only through better wages but also through better non-monetary conditions. The last point goes against the idea that higher wage jobs are not more desirable because they are viewed as compensation for worse working conditions.

Finally, we take a first step in assessing whether the magnitudes of our empirical findings could reasonably be accounted for in a model of directed search. We argued that both a positive wage elasticity as well as applications only to the low but not to the high wage are qualitatively consistent

[6] Akin to our study, one could take jobs without wage announcements and experimentally vary verbal indicators. This goes beyond this initial study presented here.

[7] We find empirically that a more recent posting date is associated with significantly more job applications even if the delay is only a day. While this feature is not present in most models we are aware of, it is broadly in line with the idea that that job seekers fear that older vacancies might have already been filled: Albrecht et al. (2020) 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.

[8] We did have to vary the postcode and the date of posting to prevent suspiciously similar vacancies, but these variations were set independently of the wage. We carefully investigate these dimensions in the empirical analysis.

with directed search. Here we use the simplest model from the literature to check whether it has the potential to generate quantitative predictions in the ball-park of our empirical estimates. The only bells and whistles that we add is to allow workers to apply to multiple job as in Albrecht et al. (2006). This addition is necessary as otherwise those who apply to a low wage vacancy have no further applications left with which to apply to a higher wage. Still this basic model has predictions that are far off our empirical targets: the implied wage elasticity of job applications in the model is an order of magnitude higher than its empirical counterpart, and those who apply to a low wage would with certainty also apply to the high wage.[9]

While our model omits a number of features that we discuss extensively in light of this failure of the most basic model (e.g., wage progression, wage dispersion, lack of commitment to the wage announcement, partially directed search, inattention to the precise wage offers), the over-reaction of the queue length is not very surprising given recent findings on the inability of such models to sustain much wage dispersion. Hornstein et al. (2011) show this failure both for random as well as directed search models, and it arises in directed search because even tiny deviations downward from the average wage reduce the queue length so substantially that it falls to zero so that no-one would apply any more. In our setting we increase the wage, and again the queue length reacts too strongly but in the upward direction. Hornstein et al. (2011) investigate a large number of extensions to their main (random) search model to rectify the problem, but only the introduction of on-the-job search is quantitatively promising. Therefore, we chose to explore this extension quantitatively in our directed search setting.

On-the-job search is promising because it implies that job spells are far shorter than employment spells, and differences between offered wages are less important as better options arrive while working. Multiple job applications have not been combined with on-the-job search, but we show that this can be done in a tractable way by allowing for competition among employers as in Postel-Vinay and Robin (2002). While we purposefully keep the model simple so as not to add too many features, this adjustment alone moves the implications on the wage elasticity and the probability of applying to the low but not the high wage very close to our empirical counterparts. We discuss the technical details that make it challenging to study the second implication in richer models with more heterogeneity or richer wage contracts, and render this beyond the scope of this paper. Yet, the fact that this stylized extension already matches the empirical observations closely suggests that richer extensions that keep this feature have a similar chance of success.

While each piece of evidence has its own challenges, in sum they suggest that the directed/competitive search approach that is becoming increasingly popular in the labor-macro literature might be appropriate to capture the application elasticity of workers within the segment of the vacancy data that posts wages. To present our approach concisely, 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 on estimating the wage elasticity. The three secondary findings are split into empirical ones relating to

[9] We recall here that for ethical reasons we only observe how workers save jobs in order to apply later, rather than their actual applications. When we interpret our findings we assume that the elasticity of saves corresponds to the elasticity of applications. That is, we assume that saves lead to applications in a constant ratio independent of the offered wage. For real vacancies we provide supportive evidence that this assumption seems justified.

the reservation wage property and the perceptions survey in section 5, while the calibrated directed search model is relegated to a separate section 6. Section 7 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 discussed above, the idea that workers target their job applications strategically to particular wages is deeply rooted in the literature on competitive and directed search surveyed in Wright et al. (2019). In the macro literature, Lentz and Moen (2017) consider future job-to-job mobility of differently paid workers to see if they target different wages. Braun et al. (2016) consider the mobility out of different levels of unemployment benefits. Engelhardt and Rupert (2017) calibrate a model that accommodates both directed and undirected search. Support for directed search is generally mixed. None of these actually considers application behavior because it is hard to observe in most datasets, but rather focus on hires.

While the idea to use actual application behavior has been around, the number of studies investigating this is limited due to scarce data on application behavior or other measures of interest in a job. Existing studies using observational data usually observe different vacancies with different wages. As mentioned in the introduction, the key challenge is to ensure that these are identical except for the wage. Otherwise one obtains counter-intuitive outcomes: Faberman and Menzio (2018) 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 (2020) 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 title, the relationship reverses: higher wages attract more applicants. Even job titles do not seem to be sufficient to make job ads comparable: Banfi and Villena-Roldan (2019) 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 when job titles are identical. Therefore, we propose our alternative strategy that relies on explicit randomization.

Our study obviously also relates to the larger monopsony literature on wage setting, surveyed, e.g., in Boal and Ransom (1997), Manning (2011) and Manning (2021), who also review search models that can explain the number of hires and the overall workforce of the firm in light of wage changes. Our setup is most closely comparable to the extensive margin labor supply elasticity with respect to the wage change of a particular firm (as opposed to changes of all firms that entails also all general equilibrium effects). Empirically it has been challenging to estimate such elasticities, as one needs credible exogenous sources of variation in wages. A few studies exploit quasi-experimental variations in wages in public-sector occupations (Staiger et al. (2010) and Falch (2010)) and find relatively low elasticities (all between 0 and 5). These studies do not directly estimate a labor supply elasticity but either a recruitment or quit rate elasticity. As Manning (2021) points out: a link to the employment elasticity is then made using the result from Manning (2003) and Manning (2011) that the elasticity of the recruitment rate should, on average, be equal to minus the elasticity of the quit rate because, for job-to-job moves, one firm's quit is another firm's recruit. The consensus from this literature is that the labor supply elasticities faced by firms are low, and are therefore indicative of substantial market power on the firms' side. Our focus here is explicitly on the job application margin which is at the centre of directed search models rather than as the reaction in total new hires which is the focus in the monopsony literature.

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 their 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. 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.[10] 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 (2020) for the US once they control for job title. This suggests an exciting possibility: our magnitude suggests that controlling for job titles in large datasets as in Marinescu and Wolthoff (2020) might indeed isolate rather similar vacancies and provide sensible magnitudes on the application elasticity, corroborating its potential as a way to use large new datasets that unfortunately do not offer natural experimental variation.

Studies with temporary jobs are likely to have different elasticities, for many reasons including possibly a different ratio of job benefits to search costs. Abebe et al. (2021) advertise clerical positions in Ethiopia and find qualitatively similar results to Dal Bó et al. (2013), though the wage elasticity of assessment centre attendance is somewhat lower at 0.45. This is likely due to the different market segment they study: in their case the job is temporary and scheduled to end after three months, and to the extent that application costs are independent of job duration one might expect a higher elasticity for long-term jobs. They also document that application costs are very high and are driven mostly by mobility costs that are unlikely to depend on job duration, and they show benefits of an application subsidies. Dube et al. (2020) analyze observational data from Amazon Mechanical Turk and revisit five field experiments with randomized wages on that platform, and find a very low elasticity of around 0.1

[10]The fact that higher wage offers attract higher quality applicants is confirmed by Ashraf et al. (2020) in a field experiment on the recruitment of nurses in Zambia. However, they also find that higher wages attract less prosocial applicants. A similar result is found by Deserranno (2019), who shows that higher wages are perceived by applicants to entail a lower positive externality to the community.

to accept new or additional jobs on that platform. These experiments consider extremely short-term tasks at piece-rates in the order of a few cents and therefore focus on a different market than full-time long-term jobs.

In a different context and with a different research focus, Leibbrandt and List (2015) post (real) vacancies for administrative assistants 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 participation and negotiation behavior. Mas and Pallais (2017) 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, 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.[11]

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.[12] Nevertheless, the database is regarded as reliable by the UK government, who strongly encouraged 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.

Our study also relates to lab experimental work with students that has investigated whether they do play according to predictions of directed search, with surprisingly high amounts of success (Cason and Noussair (2007); Anbarci and Feltovich (2013); Helland et al. (2017); Kloosterman (2016)). Those studies have been rather stylized, with no intention to replicate a real labor search environment. Our setup uses real job seekers who spend their time searching through mostly real jobs, and is therefore more akin to a field-in-the-lab setting which we had to set up to replicate the field as ethical considerations did not easily allow us to collect consent for our study in the field directly.

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

[11] There 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.

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

as whether minorities react differently to higher wages or employer ethnicity or gender.

3 Experimental Design

The setup, recruitment process and institutional setting are described in detail in Belot et al. (2019). The latter paper evaluates the effects of an intervention providing tailored advice. This experimental intervention 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 websites 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.[13] 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 (2017), Marinescu and Wolthoff (2020) and Marinescu and Rathelot (2018). 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. (2019) - 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.[14]

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

[13] Based 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

[14] In an exit survey the vast majority of participants (86%) said that this did not affect their search behavior. This is likely due to the very low numbers of artificial vacancies and to the fact that fake advertisements are common in any case on online job search sites (see footnote 12 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

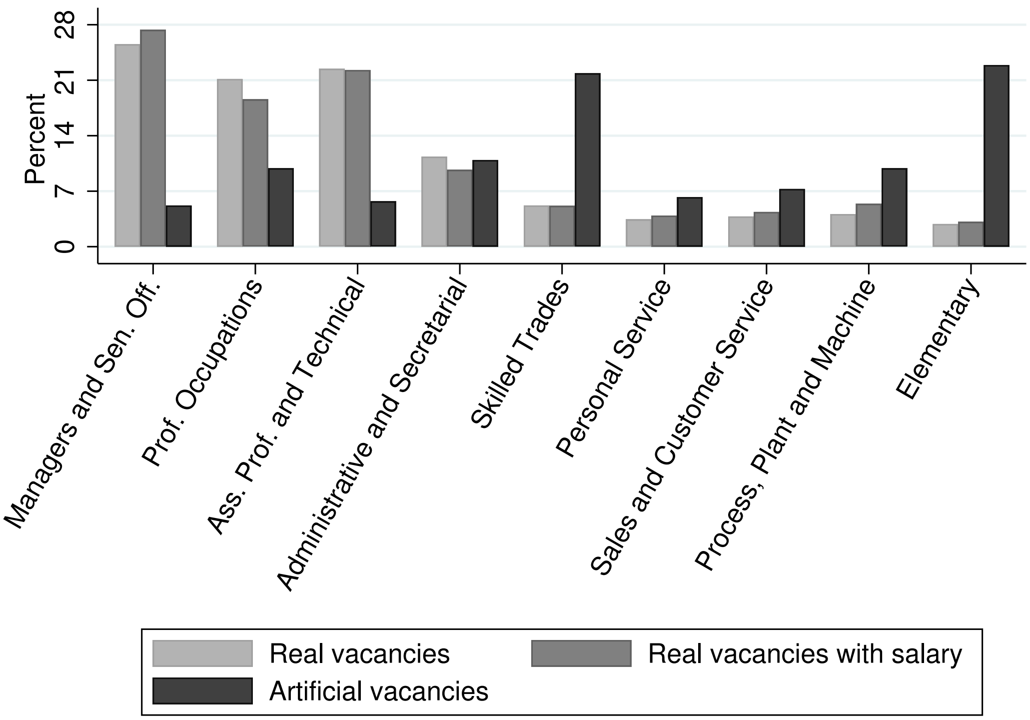

never exceeded 2%. We also checked whether the share of artificial vacancies within all vacancies saved by participants did not exceed 2%, and adjusted the number in subsequent weeks in case it did. The vacancies were added to the database of real vacancies during the days on which lab sessions for participants took place. Each artificial vacancy was only active during sessions of a particular week, such that participants would never observe them in multiple sessions.[15] In this section we describe the 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

[15]To be precise, all job search sessions in our study occurred on Mondays and Tuesdays (see section 3.3), while the artificial vacancies were posted on the Fridays, Saturdays or Sundays prior to the job search sessions.

vacancy text etc.). From each selected vacancy we removed all identifying information (company name, contact person, telephone number, website, etc.).[16] 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 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.[17] The vacancies are classified by the first digit of their UK SOC code. We present the same distribution for the selection of vacancies that post a salary in the middle bars of Figure 1. The distribution of vacancies with posted salaries is quite similar to the overall distribution. 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.[18] 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 minimum to a maximum.[19] 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

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

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

[18] Note 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.

[19] One might wonder whether the posted wage is indeed the wage that employers eventually pay to the hired worker. While it is difficult to obtain representative data on this, we obtained some evidence that this is indeed the case. In a brief survey among employers that posted vacancies, we find that more than 50% of the employers pay exactly the wage posted in the ad (or the midpoint of the posted wage range). Given the low response rate (<8%) in the survey we do not consider the survey any further in this study, though brief results can be found in the Online Appendix (tables 18 and 19). We thank an anonymous referee for suggesting this question.

Figure 2: Distribution of posted salaries

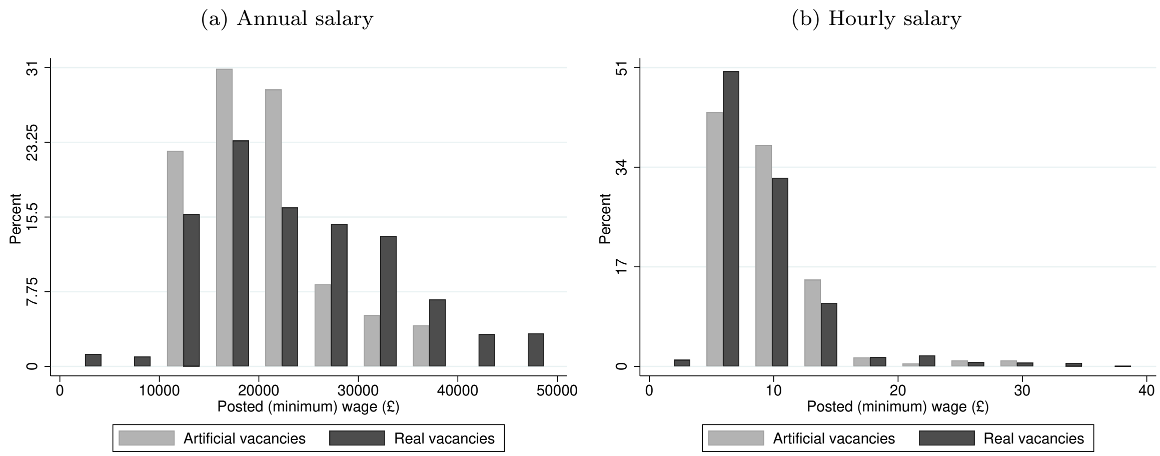

mass in the tails of the distribution, but other than that they are quite similar.[20]

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 resumes 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 most 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.[21] 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 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 posting dates of the vacancies were set to always precede the job search sessions by 1 - 3 days, making the vacancies more prominent within the search results of job seekers (see also footnote 15). Within the pair, the posting date typically varied by 1 or 2 days, to make identification of the pair less likely.

The construction of the artificial vacancy pairs was done in the following manner. We created

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

[21] Some audit studies do send pairs of resumes to the same vacancy, such as Farber et al. (2016).

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%.[22,23] 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 difference would be. We made sure however, that in case of a salary reduction the new salary would not be below the minimum wage.[24] As a result our sample contains relatively more pairs with a positive wage difference among low wage vacancies than among high wage vacancies, while also overall we have more pairs with positive wage differences than with negative wage difference. 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.[25] In 32 pairs there is negative wage difference, in 75 pairs there is 20% positive wage difference and in 54 pairs there is 40% positive wage difference.[26] For vacancies that originally offered a wage range, we multiplied both their minimum and their maximum wage by the same factor.[27]

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 of posted vacancies at the finest (3-digit) occupational code level. We compute for each 3 digit code the magnitude of a 20 or 40 % wage difference in terms of the standard deviation of wages within this occupation.[28] We find that on average a 20% higher (or lower) wage corresponds to 0.44 of a standard deviation, while a 40% higher or lower wage 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% wage difference is not likely to be outside the support of

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

[23]One 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.

[24]In case the assigned negative wage difference resulted in a wage below the minimum wage, we assigned a (random) positive wage difference 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.

[25]In 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 20, 40 or 60% higher. 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.

[26]Initially we only created artificial pairs with positive wage difference, while later we decided to also add some with negative wage differences. In addition, not all negative wage differences were feasible due to the minimum wage lower bound. As a result we have many more positive than negative wage differences.

[27] For example, a vacancy originally offering a wage between 10 - 12 pounds, would be manipulated to offer 12 - 14.40 pounds, such that both the minimum, the maximum and the mean of the range are 20% higher.

[28] These 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. Because our artificial vacancies all lack a company name, the higher wage might look particularly suspicious. This is unlikely however, because we find that vacancies that lack a company name actually post on average a 13% higher wage, conditional on the 3-digit occupation. We also reproduce Figure 3 including only real vacancies without company names and find very similar results (see Figure 6 in the Online Appendix).

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

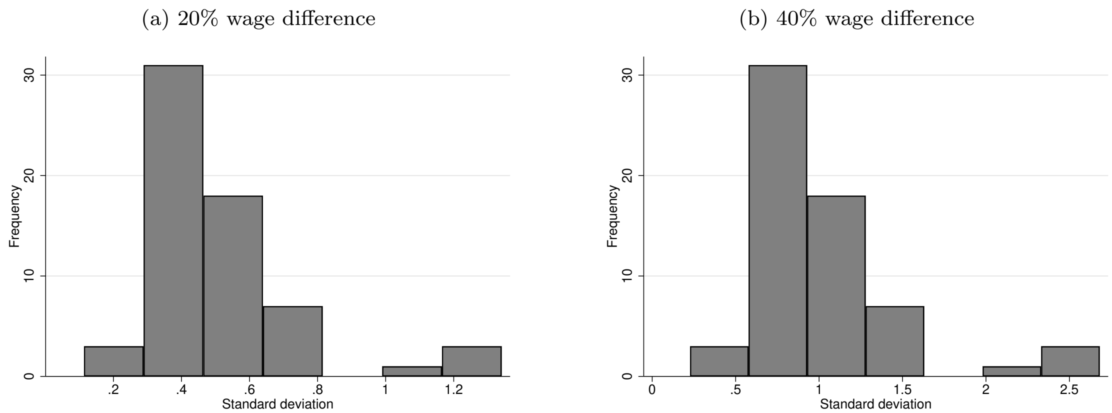

the wage distribution.

At the end of the study we performed a small survey to assess whether participants felt that the artificial vacancies had affected their behavior. When asked whether they were able to distinguish the artificial vacancies from real vacancies, 68% answered 'never' or 'rarely'. 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 A.3 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

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.

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.

Participants were asked to search for jobs using our job search engine (described later in this section) for a minimum of 30 minutes.[29] 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 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). We handed out instructions on how to use the interface, which we also read aloud (the instructions can be found in the Online appendix). 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.[30] 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

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

[30] We 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.

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. (2019). The key goal of the alternative interface was to offer suggestions to job seekers about occupations that might be of interest to them. This was achieved by 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: wage elasticity

Our empirical analysis focuses on the question whether, all else equal, higher wages increase job seekers' interest in vacancies. To answer this question we estimate the wage elasticity, which is the percentage change in interest for a 1% change in wage. We first discuss the measures of interest that we observe in the experiment, before presenting elasticity estimates based on the real vacancies and experimental (artificial) vacancies.

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.[31] 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 7 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 Results: observational data

We first explore how wages are related to job seeker interest amongst *real* vacancies that were posted during our study. We observe how often each of these vacancies was saved or viewed by participants in the study. The posted wage is highly correlated with other characteristics of the vacancy and thus the estimates reflect associations rather than causal relationships. We investigate to what extent these associations remain when controlling for an increasing number of vacancy characteristics, but clearly information conveyed through the flow text cannot be controlled for. The results are intended as a benchmark to compare to both previous literature as well as our experimental results.

Since the number of saves (S_i) is a count variable, we perform a Poisson regression on the logarithm of the offered wage, controlling for vacancy characteristics X_i:

$$S_i = \exp(\alpha + \beta \log(w_i) + \delta X_i + \epsilon_i)$$

The parameter of interest is β, which captures the wage elasticity. Our sample contains vacancies that were posted in the Edinburgh area during the experiment, in order to make the analysis comparable to our experimental results. From this set we select all vacancies that post a wage (46%) and focus on those that post the wage in *annual* terms (54%). The last restriction is used to prevent a problematic comparison of hourly, daily, weekly, monthly and annual wage announcements, although we reproduce the main regression with vacancies announcing hourly wages as a robustness check.

Estimation results are presented in Table 2. We find in column (1) that a higher wage is associated with significantly less saves, not controlling for any vacancy characteristics. The elasticity is -0.73 and highly statistically significant. The association remains (though with slightly smaller magnitude) when controlling for four-digit 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 mostly statistically significant, but only marginally change the negative wage coefficient (columns (3)-(6)),

[31] We have also checked whether the probability to apply (conditional on saving) depends on the wage of the job, and find that no such relation exists once we control for occupation and individual fixed effects. Thus, the saving decision can be considered a valid proxy for applying, even once we differentiate the posted wage in our artificial vacancies.

Table 2: Effect of wage on number of saves: real vacancies

			Poisson regression				Log-log regression
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log(Salary)	-0.73***	-0.69***	-0.61***	-0.61***	-0.62***	-0.19**	-0.44***
	(0.0093)	(0.032)	(0.038)	(0.038)	(0.040)	(0.078)	(0.059)
Temporary contract			-0.18	-0.18	-0.15	-0.32***	-0.19**
			(0.11)	(0.11)	(0.11)	(0.067)	(0.075)
Part time			0.71***	0.71***	0.68***	0.21*	0.55***
			(0.23)	(0.23)	(0.23)	(0.13)	(0.096)
No company name			-0.51***	-0.55***	-0.53***	-0.11	-0.37***
			(0.057)	(0.054)	(0.056)	(0.079)	(0.072)
No contacts in ad				0.56***	0.61***	0.50***	0.24***
				(0.20)	(0.20)	(0.13)	(0.088)
Constant	284367.0***						2.35***
	(96459.4)						(0.54)
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	7165	6804	6731	6731	6731	4898	7091

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

which remains close to -0.6. When analyzing jobs that post an hourly wage, we find a similar negative wage coefficient that is smaller in magnitude (-0.19, column (6)). Finally, also a log-log regression leads to very similar results (column (7)). When using viewing a vacancy as the outcome (instead of saving a vacancy) we find very similar results, presented in Table 20 in the appendix.[32]

Thus, even when comparing jobs with the same (narrowly defined) occupation, similar hours and similar type of contract, a higher wage is associated with *less* job seeker interest. This result is in line with the findings of Marinescu and Wolthoff (2020) who analyse jobs from Careerbuilder.com (when not controlling for job titles) as well as the results of Faberman and Menzio (2018). We expect the negative association to be driven by job characteristics such as required qualifications or experience that we cannot control for and that are likely positively correlated with the wage.

4.3 Results: experimental data

We now consider our *experimental* vacancies, which were created in pairs in which one kept the original wage and the other was randomly assigned a higher or lower wage. A simple within-pair comparison yields 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, in contrast to the findings from observational data, we find that for randomly assigned wages, a *higher* wage attracts *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 differences and estimate elasticities, 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:

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

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 difference 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 altered wage.[33] To correct for the correlation between these four vacancies, we cluster standard errors at the pair level.[34] As an

[32] In addition we reproduce Table 2 using only real vacancies that were posted in the three days prior to each of our experimental sessions. Doing so, we ensure that we consider real vacancies that are likely to be actually considered by the job seekers. Results are very similar, see Table 21 in the Online Appendix.

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

[34] The 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.

Table 3: Effect of wage difference on number of saves/views

	Saves			Views		
	(1)	(2)	(3)	(4)	(5)	(6)
Salary difference (in %)	0.70**	0.69*	0.92***	0.70**	0.71**	0.86***
	(0.44)	(0.45)	(0.43)	(0.35)	(0.36)	(0.29)
Appearing first			0.58***			0.50***
			(0.13)			(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).

extension to this simple specification we add additional controls for the geographical location of the job and for whether the vacancy had the more recent posting date (relative to its counterpart) and thus appeared first within the search results.

Before turning to the results we briefly motivate our choice to perform the analysis at the vacancy-level. Alternatively we might consider individual job seeker decisions as the outcome (that is, whether individual i saves vacancy j). We prefer an analysis at the vacancy-level, because we are interested only in variables that vary by vacancy (most importantly the wage, but we'll also include other vacancy characteristics). Since the wage varies randomly (within pairs), there can be no correlation with individual factors and we see no gains to performing analysis at the individual level. Furthermore, a linear probability model or a logit model for individual choices introduces other challenges: neither estimates a constant wage elasticity, which is the key parameter that we are interested in. Finally, an individual choice model would imply using a sample containing almost exclusively zeros as the outcome: the probability that one particular job seeker saves one particular vacancy is very close to zero.

Estimation results are presented in Table 3. We report $exp(\beta) - 1$, which is the percentage change in saves due to a 1-percent difference in the wage (the elasticity).[35] 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 saves. The posting dates were assigned independent of the wage, and as expected we find that the wage coefficient only changes slightly (0.92). This indicates that our elasticity measure is relatively robust and constitutes our main finding for this section.[36]

[35] Note that the number of observations reflects that in Poisson regression all pairs for which none of the vacancies were ever saved are dropped.

[36] As an alternative to the Poisson model, we can estimate a log-linear model where the dependent variable is the

Our elasticity estimates are very similar in magnitude to the results of Dal Bó et al. (2013), who report that a 33 % higher wage 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 Marinescu and Wolthoff (2020), 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. (2021) 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. However, they consider three-month temporary jobs and a similar wage difference for permanent jobs would lead to larger reactions if job seekers trade off fixed application costs with the net present value of wages on the job.

Rather than using saving a vacancy as the outcome variable, we can also consider the decision to view a vacancy. The results are presented in columns (4)-(6) in Table 3 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.

4.3.1 Wage increases and wage decreases

The wage elasticity might vary across different kinds of vacancies. First, we consider differences between positive and negative wage differences and second we consider differences by baseline wage level of the vacancy pair (low or high).[37] Table 4 reports elasticities of saving and viewing, distinguishing between positive and negative wage differences. We find that the elasticity of the number of saves is significant for the positive wage differences where we have the majority of observations (129 pairs), and shows no statistical significance for our limited observations with negative wage differences (32 pairs). The elasticity of the number of views is significantly positive for types of wage differences.

Negative wage differences raise the problem that they might lie below the minimum wage. So any negative wage difference 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 positive wage differences were also possible at low wage vacancies where negative wage differences were not feasible (which constitutes 40% of our sample). When we separately evaluate the effect for low wage vacancies with positive wage difference, high wage vacancies with positive wage differences, and high wage vacancies with negative wage difference (column 2) the effects are concentrated in the first set.[38] 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

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. Finally, we show in Table 22 in the appendix that the results are almost identical when also Glasgow pairs are included in the analysis.

[37] Note that while we have no a priori reason to expect different responses to increases and decreases, our model calibration in section 6 does show slightly different elasticities with respect to wage increases and decreases.

[38] Rather 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*).

Table 4: Effect of wage difference on number of saves/views: positive and negative wage differences

	Saves		Views	
	(1)	(2)	(3)	(4)
Salary dif.*positive	0.96***		0.82***	
	(0.45)		(0.29)	
Salary dif.*negative	0.24		1.84*	
	(0.74)		(1.62)	
Sal. dif.*low*positive		1.21***		0.88***
		(0.61)		(0.35)
Sal. dif.*high*positive		0.31		0.63*
		(0.46)		(0.47)
Sal. dif.*high*negative		0.23		1.84*
		(0.74)		(1.62)
Appearing first	0.57***	0.56***	0.50***	0.50***
	(0.13)	(0.13)	(0.076)	(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).

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.

5 Secondary empirical findings related to directed search

The leading quote for this article distinguishes between a rather passive role of firms in attracting job applications, and a more proactive role in which the wage is used competitively to attract applicants. The latter is usually captured in models of directed search. Our main result from the previous section - the significant positive elasticity of job applications with respect to the experimental wage variation across a large number of occupations - supports this view. While this is arguably the most well-known implication of directed search and a main focus of our analysis, we now investigate further properties of such models and whether they find support in our setting. We provide results on three dimensions. First, we consider whether a reservation wage property holds in our data. That is, do job seekers that save (or view) the lower wage vacancy in pair, also save the higher-wage vacancy in a pair? Directed search models do not have this implication, but one can construct plausible alternatives that do, as discussed below. Second, we present results from a perceptions survey, where participants were asked how they perceived competition, required quality and non-monetary working conditions for each of our artificial vacancies. These perceptions provide insights on beliefs regarding the likelihood of getting the job, which play a key role in directed search. Third, we evaluate whether the magnitudes of our empirical findings can be reproduced in a calibration of a directed search economy, though this is relegated to the next section due its different methodological nature. While these investigations are more of an explorative nature, they nevertheless in sum seem to support a directed search interpretation.

5.1 Is a 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 higher wage 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. The defining feature here is that workers can cheaply send as many applications as they like, and they only refrain from applying from a job they know about if that job is so unattractive that the worker would not accept employment.

Competitive or directed search models differ in their assumption that workers can apply to strictly less vacancies than they find acceptable - 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. While directed search models differ in their exact formulation, we are not aware of any version that satisfies a reservation strategy as outlined above.[39]

Therefore we investigate whether a 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. The probability of not saving the

[39] Models in this literature tend to feature a continuum of jobs, while each worker in a given period can only send a finite number N of applications, where often this number is set to $N = 1$. These models never satisfy a reservation wage property by construction: whenever a positive fraction of jobs exceed the reservation threshold workers cannot possibly apply to all of them. This inability to apply to a large number of options is exactly what induces workers to be selective. Note that we use the term "reservation wage strategy" to mean that workers apply to all jobs they encounter above a reservation threshold. Other papers use this term differently. For example, Peters (2010) derives a rule in a directed search model where workers apply uniformly to all wages above a threshold. This would not be a "reservation wage strategy" in the terminology of this paper as for any threshold there are many jobs with wages that are above it to which workers do not apply.

Table 5: Outcome variable: Probability of not saving high-wage, conditional on saving low-wage

	Saves			Views		
	(1)	(2)	(3)	(4)	(5)	(6)
High appears first			-0.29***			-0.27***
			(0.053)			(0.046)
Constant	0.45***	0.42***	0.53***	0.40***	0.38***	0.49***
	(0.030)	(0.030)	(0.037)	(0.028)	(0.028)	(0.034)
Sample	All	List both	List both	All	List both	List both
N	278	267	267	520	500	500

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

high-wage vacancy, conditional on saving the low wage vacancy and vice versa are denoted by:

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

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

where $S_l = 1$ if the low wage vacancy in a pair was saved, and similar for S_h. We present probability (2) in the first column of Table 5. Note that for ease of representation and comparison, the first column presents an estimate from regressing the conditional probability (equation 2) only on a constant. Thus, the estimate 0.45 means that in 45% of the cases where someone saves a lower-wage vacancy, they do *not* also save the higher wage vacancy from the pair. The estimate is strongly significantly different from zero, which would be the null hypothesis in the case that the reservation wage property holds.

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.[40] The advantage of our experimental setup is that we can observe which vacancies a job seeker has 'listed' on their screen (meaning which vacancies were certainly seen on the screen by the job seeker). So we condition on cases where both vacancies were listed by the job seeker. As is shown in column (2), this hardly changes the results, and still in 42% of the cases where someone saves a low-wage vacancy, they do no also save the high wage vacancy.

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. One 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) infer that more than a third of vacancies on a US search platform are already taken, and Albrecht et al. (2020) 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

[40]If, 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.

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.[41] 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 control for whether the high wage vacancy appeared first (relative to its low wage counterpart) in column (3). In line with expectation, this reduces significantly the probability of not saving the high wage vacancy. The estimates imply that $P(S_h = 0|S_l = 1)$ is 0.53 when the low-wage job appears first. It is 0.29 lower (i.e., equal to 0.24) when the high-wage job appears first. Thus, even when the high wage job appears first, almost a quarter of people who save the low wage job, do not save the high wage job. 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 24% and 53% that we assess here. This was not feasible in our setting as this would have appeared too unrealistic, and absence better information we use the average estimated effect of 42% as our benchmark.

We can compute similar probabilities for viewing vacancies. These results are reported in columns (4)-(6) of Table 5. 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 0.22. All of these are very close to the results from saving the vacancy, 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.

We present the reverse probability (equation 3) of not saving a low-wage vacancy conditional on saving the high-wage vacancy in Table 12 in the Appendix. This probability equals 0.52, which is, as expected, slightly larger than the probability in Table 5. Nevertheless, the difference between the two is smaller than one might have expected.

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 relevant, and we have collected several individual characteristics, our study has not been designed to answer these questions. We restricted ourselves to no more than 2% research vacancies in the sample of real vacancies to keep the environment realistic. Also the number of job seekers is relatively limited, such that we lack power to delve into questions of heterogeneity. Keeping this in mind, we have performed several analyses concerning the relation between the response to the posted wage and individual characteristics. We found no evidence of heterogeneous responses in terms of gender, ethnicity (share of whites), age, education (share with university degree), family type (share with children) unemployment duration (share with less than 6 months unemployment duration at the start of the study), time or risk preferences. However, 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.

[41] An alternative explanation could be stock-flow matching. In our setting this seems less relevant as most job seekers are well beyond their first weeks of unemployment.

We now shortly discuss six alternative reasons why someone might only save the low-wage vacancy, and show that none of them can convincingly explain our findings. A more detailed discussion can be found in appendix section A.7, including results in Tables 13-16.

First, 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. 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). Reproducing our results on the remaining observations provides very similar numbers and thus we conclude that identifying pairs of vacancies does not drive the results.

Second, we further consider whether individuals are able to identify artificial vacancies, by selecting participants that reported in the exit survey that they could 'Never' (48%) or 'Rarely' (20%) identify whether a vacancy was not real. This implies dropping 654 vacancies viewings (57 % of the observations) or 347 vacancy saves (56 % of the observations). We find a slightly lower value for $P(S_h = 0|S_l = 1) = 0.36$, while the importance of the posting order is a bit larger. As a result, for cases where the high wage appeared first, the probability of saving the low wage but not the high wage is reduced to 0.09. However this probability is based on a very small sample of 34 vacancy observations, implying that a 95 % confidence interval includes values between 0.03 and 0.23 and thus the number is not significantly different from our previous estimates, while it is significantly different from zero. Note that this sample is small and selective because only participants present in the final session completed the survey. In particular, those that apply more often to higher wages than lower ones are less likely to succeed if such jobs are harder to get, and are thus more likely to end up in this survey.[42]

Third, the slight variation in geographical location within a pair might affect the saving decision (although all vacancies were located within the Edinburgh area). Through computing the geographical (straight-line) distance between the postcodes, we select only pairs for which the distance is at most one kilometer (approximately 0.6 miles, comprising roughly 50% of the pairs). Using this subsample for which geographical variation is unlikely to affect choices, we find almost identical probabilities and thus conclude that location differences do not drive our results.

Fourth, we focus only on individuals that *list* relatively few 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 probabilities for the subsample of individuals that lists less than the median number of vacancies per week (on average). When both vacancies have been listed this probability is 0.30, while it was 0.42 in the complete sample. This is a natural result that also follows from the directed search model that we discuss and calibrate in section 6: the smaller the set of potentially interesting vacancies, the larger the probability of saving each particular vacancy. In the appendix we discuss that this also holds for the reverse probability (saving high but not low). In conclusion, it appears that even in this sample a significant share of the

[42] Only about 50 % of the participants completed the exit survey. Much of the reduction in the statistic can be attributed to this particular selection. Among *all* participants that completed the exit survey after three months of job search, the probability is 0.16.

individuals prefers the lower wage vacancy.

Fifth, we select only vacancy pairs with small wage differences to eliminate the concern that large salary differences imply a wage that might look suspiciously high (or low) and thereby affect saving and viewing choices. Selecting only pairs with a 20% salary difference, we find a very similar probability of not saving the high wage vacancy while saving the low wage vacancy (0.41 vs 0.42 in the baseline).

Last, we discuss a very different concern that might arise: do participants learn and become more aware of artificial vacancies over the 12 weeks, or become better at identifying the artificial vacancies? To investigate whether learning occurs, we split the observations into three 4-week periods and compute the ratio of savings per viewing for artificial vacancies. If participants get better at identifying artificial vacancies (once they view one) this ratio would be expected to decrease over time. We find that it is 0.53 in weeks 1-4, 0.51 in weeks 5-8 and 0.58 in weeks 9-12. Two-sided t-tests for equality show that none of these ratios differ significantly from each other. In addition, 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 any differences in the saving rate to differences in the pool of artificial vacancies and differences in the pool of participants (due to attrition).

5.2 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 there are differences in expected (unobserved) characteristics in other dimensions. A higher wage could for example be interpreted as a signal of worse working conditions. If job seekers assume that higher wages are only offered to compensate for differences in working conditions, it is rational (for some) to prefer the low wage job.

To investigate whether this is the case, we designed and conducted an online survey. Participants were non-students, recruited through the behavioural lab of the University of Edinburgh, and received a £20 voucher. 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 is on average 31 years old (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%).[43]

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

[43] Ideally one would have collected 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.

Table 6: 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

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 those theories of compensating differentials where higher wage jobs are not more desirable because they simply compensate for worse working conditions (see Rosen (1986)).

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. 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.[44] 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.[45]

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 6. 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)

[44] 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.

[45] To assess the quality of the measurements of perceptions about the vacancies, we compute Cronbach's alpha for each set of individuals that judge the same 20 vacancies. There are 16 of such sets. The distribution of the 16 values is presented in the appendix in panel (a) of Figure 10. 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 10), we find that agreement is particularly high for question 1 and 3, while somewhat lower for question 2.

Table 7: Effect of salary on saves/views, controlling for perceptions

	Saves		Views	
	(1)	(2)	(3)	(4)
Salary difference from original salary	0.92***	0.74*	0.86***	0.74**
	(0.43)	(0.58)	(0.29)	(0.37)
Q1 (quality) standardized		-0.053		0.13
		(0.18)		(0.17)
Q2 (competition) standardized		-0.087		-0.16
		(0.15)		(0.098)
Q3 (working conditions) standardized		0.35*		0.30**
		(0.22)		(0.14)
Appearing first	0.58***	0.59***	0.50***	0.50***
	(0.13)	(0.13)	(0.075)	(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).

and will thus be harder to get. These results also reject the alternative hypothesis that high wages compensate for worse working conditions (i.e., act simply as 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 positive wage elasticity found in section 4.3 changes if we control for perceptions about the vacancies. To do so, we compute the average responses across participants for each vacancy, using linear values ranging 1-5 for the choices (where a higher number means a "higher" response). 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, in order to correct for heterogeneity across participants. The mean standardized responses are then included as control variables when regressing log-saves on the salary difference with a specification that is otherwise identical to the one used in Section 4.3. 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 7. Column (1) reproduces the main finding from Table 3. 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,

Table 8: 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	0.083	0.47*	0.43
	(0.31)	(0.42)	(0.25)	(0.30)
Difference Q1 (Quality)		-0.040		-0.028
		(0.10)		(0.067)
Difference Q2 (Competition)		-0.066		-0.078*
		(0.11)		(0.047)
Difference Q3 (Working conditions)		0.066		0.13*
		(0.10)		(0.076)
High appears first	0.29***	0.29***	0.24***	0.25***
	(0.076)	(0.076)	(0.050)	(0.050)
Constant	0.45***	0.44***	0.36***	0.38***
	(0.090)	(0.10)	(0.073)	(0.080)
P-val joint sign. Q1-Q3		.88		.22
Individual fe	yes	yes	yes	yes
N	278	278	520	520

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. 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.

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 23 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 7 we find that results are very similar when considering views instead of saves.

An alternative is to look at the conditional probabilities (analyzed in section 5.1). 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 variables. 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 within pair differences in wages and differences in perceptions as regressors. The regression specification is:

$$\mathbb{1}_{(S_h)_{pj}} = \alpha + \beta \Delta w_p + \sum_{k=1}^{3} \gamma_p \Delta P_{k,p} + \delta PO_p + \pi_j + \varepsilon_{pj} \quad (4)$$

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 between the high and low wage vacancy in pair p are given by $\Delta P_{k,p}$. We also control for the posting order: a dummy for the high wage 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 8. 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 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 statistical significance, 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 and working conditions are detected as significant factors.

6 Calibrating a directed search model

In the empirical section we presented our main finding:

1. Jobs with higher wages attract more interest from job seekers, with an elasticity between 0.7 and 0.9 (Table 3).

We also documented as additional insight:

2. Some people apply only to the lower wage job within a pair, and external evaluators perceive high-wage jobs as harder to get. In particular, of those who apply to the low wage, roughly 40% do not apply to the high wage even if both are displayed on their screen (Table 5).

We discussed that these findings are qualitatively in line with predictions from directed search, which might lend itself therefore as a natural candidate to account for this. It remains an open question whether the magnitudes that we find are plausible as well, and we investigate this as a third additional step to understand whether directed search might be a plausible way to account for the empirical evidence. 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 points. An obvious candidate is the canonical directed search model with homogeneous workers and firms, where workers can apply to one job per period, and in equilibrium all firms offer the same wage (as in Burdett et al. (2001) or 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 tweak it to a minimal degree to allow us to investigate the second of our two empirical findings.

In order to assess the second finding we need a model with multiple job applications, otherwise it follows trivially that a worker who applies to the low wage will not apply to the high wage since (s)he has no further applications left.[46] 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 job seekers who can see all of the (continuum of) vacancies, this feature is unattractive for our present purpose: if we fix a low and a high wage vacancy, and a job seeker applies to the low wage, his chance of also applying to the high wage is again zero. Even though he now has additional applications to send, given the large set of options any particular one of them has negligible chance of being selected. This changes if job seekers only see a finite number of vacancies: conditional on seeing both the low and the high wage vacancy and applying low, there is a well-defined and possibly positive probability of applying high.

Luckily, it is known since Acemoglu and Shimer (1999) that the original "see-all" equilibrium remains an equilibrium when workers see only a finite number of randomly selected vacancies.[47] This is the case if

i. workers see strictly more vacancies than the number of applications they send, and

ii. there is no wage dispersion in the original "see-all" equilibrium.

The first point is obviously necessary: if workers have weakly more applications than the number of vacancies they observe, they apply to all of them instead of being selective about where exactly they send them. The second point is also easy to understand: assume there is wage dispersion and the optimal strategy is to send one application to a low wage and one application to a high wage. If job seekers only see a finite set of randomly selected vacancies, some will only see low wages, and so they cannot implement their optimal strategy. While both ingredients are clearly necessary, they are also sufficient as we discuss in more detail below.

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 - avoids wage dispersion on the equilibrium path. So we build on it and allow workers to see a finite number of vacancies that exceeds the number of applications that they send. We treat our positive wage difference as an off-equilibrium deviation by a single firm. We discuss below a large number of extensions that will make this model more realistic and allow it to be more in line with the data. We explore a particularly prominent one of them (on the job search) quantitatively. But we note already here that explicitly allowing for wage dispersion on the equilibrium path goes beyond this paper as it would either require us to drop

[46] It might be possible to use models where workers apply once each period, but have very frequent periods: If workers apply once a day and we post our paired vacancies one day apart, then a worker could in principle apply to both. We thank a referee for pointing this out. Such models usually assume that workers know the outcome of their previous application before making the next one, which is unlikely at daily frequency that we post the two vacancies. If the worker does not know the outcome, such models resemble those with multiple applications in one period that we pursue here. Using one-period models with daily frequency would anyhow require additional adjustments, otherwise the average monthly number of applications would substantially exceed their data-counterpart.

[47] The equivalence was proven for a single job application but the logic extends directly to multiple job applications; see below.

investigating the second question or it would require a much more elaborate machinery to deal with finite visibility in such models.

6.1 The model

Assume time is discrete, β is the discount factor and each worker has has a fixed number $N \geq 1$ of job 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 $1-\kappa$ an application is not suitable. This is immediately visible to the firm, and only the fraction κ of applications is considered by firms.

It is well-known in the literature that a suitable worker who applies to a job with wage w where in expectation also $\kappa\lambda(w)$ other suitable workers queue for a job has a chance of $(1-e^{-\kappa\lambda(w)})/\kappa\lambda(w)$ to get a job offer (see, e.g., Wright et al. (2019) for a derivation). When sending the application the worker is not sure whether he will be suitable, so his chances ex ante are

$$m(\lambda(w)) = \frac{\kappa(1-e^{-\kappa\lambda(w)})}{\kappa\lambda(w)}. \tag{5}$$

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. 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.4. 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. 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

$$\begin{aligned} U(w^d) \quad &= \quad b + m(\lambda(w^d))[1-m^*]^{N-1}\beta V(w^d) \\ &\quad +(N-1)m^*[1-m(\lambda(w^d))][1-m^*]^{N-2}\beta V(w^*) \\ &\quad +[1-m(\lambda(w^d))][1-m^*]^{N-1}\beta U^* \\ &\quad + \begin{bmatrix} 1-m(\lambda(w^d))[1-m^*]^{N-1} \\ -(N-1)m^*[1-m(\lambda(w^d))][1-m^*]^{N-2} \\ -[1-m(\lambda(w^d))][1-m^*]^{N-1} \end{bmatrix} \beta V(y) \end{aligned} \tag{6}$$

where the first line reflects his unemployment benefit receipt b, as well as the probability to succeed only at the deviant firm, 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^*. Finally, the last line captures the probability of being successful at more than one firm, in which case firms bid against each other for the worker's services until the wage reaches his productivity and the continuation value is $\beta V(y)$.

A worker employed at wage w obtains

$$V(w) = w + \beta(\delta V(w) + (1-\delta)U^*), \tag{7}$$

which reflects his wage payment, and the discounted value in case he remains employed and the value in case the job breaks.

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

$$U(w^d) = U^* \tag{8}$$

since workers have to be indifferent between the equilibrium wages and the deviant wage in order for some but not all to apply there. So for any conjecture about w^* this determines the queue lengths at all other wages. Finally, the equilibrium wage is tied down by the requirement that it has to be optimal for firms. This means that it is either given by their first order condition taking into account the reaction of the queue length to different wages, or it is at the lowest bound that workers are willing to accept. For that, note that wages cannot fall below the workers' reservation wage r (defined as $V(r) = U^*$).[48] Once the equilibrium is solved for, one can use (6) 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). Locally this is given by $w^*\lambda'(w^*)/\lambda(w^*)$, though we also compute it for a large change 20% change as in our experiment.

We are also interested in the probability of applying low, but not high, conditional on observing both. As discussed earlier, calculating this probability requires setting an additional parameter X representing the number of vacancies that each job seeker observes. As long as $X > N$ it can be shown that w^* remains an equilibrium: assume firms still post equilibrium wage w^*, then the queue length at this wage is unchanged because the number of job applications has not changed. Moreover, the queue length at deviation wages has not changed locally because those workers who do see the deviation wage still have to be indifferent between applying there or at the equilibrium wage. This only changes when the wage is so high that workers' indifference requires a queue length above Xu/v, which is the upper bound when all workers who see the deviant's vacancy apply there. Clearly, if a firm did not

[48] Albrecht et al. (2006) show that this corner solution always arises at the equilibrium wage. However, when we introduce on-the-job search later on, this no longer needs to be the case.

want to deviate when the queue length was higher, it will not deviate now when it is bounded by this number, so the equilibrium wage is unchanged.

In the appendix we derive that the probability of applying to an equilibrium wage job and not to a higher deviant wage job (conditional on observing both) equals

$$\Pr(A_h = 0 | A_l = 1) = N(1-p)/(N-p) \tag{9}$$

as long as this number is positive, and zero otherwise. Here 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:

$$\lambda(w^d) = puX/v. \tag{10}$$

Recall that as long as workers apply with interior probability, they have to be indifferent between the equilibrium wage and the deviation wage, and the queue length $\lambda(w^d)$ is determined as a solution to (6) and (8) independently of exactly how many vacancies are observed. As discussed earlier, if X goes to infinity so that all vacancies are observed, the probability of applying to any vacancy including the deviant becomes zero by (10), and the probability to apply to the high wage conditionally on applying to the low wage in (9) becomes one. Only for finite X it falls to lower levels, which is exactly the reason we introduce finite observability as discussed earlier. Finally, the probability of applying high but not low can be derived analogously.

6.2 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 9. 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 5% discounting per year. Market tightness v/u follows from the UK Office for National Statistics (ONS), which, together with N, pins down the equilibrium queue length $\lambda(w^*)$.

Define the weekly job finding rate j. 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 $j = 5.2\%$ per week, which is 20.7% per month. Since the job finding rate j 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 κ from equation (5).

For the separation rate we use statistics on the UK labour market from Gomes (2012) whose estimates for the job-to-unemployment and job-to-inactivity hazard imply 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).

Finally, for the second question we need to take a stance on the number of observed vacancies X.

Table 9: Calibration of parameters

Parameter	Meaning	Value	Source
Basic model:			
Period length		Month	
N	Applications	10.4	Based on own study
β	Discount factor	0.9957	5% discounting per year
v	Vacancies	573,000	Vacancies in UK, 2013, Q4 (Office for National Statistics (2013c))
u	Unemployment	2,348,000	Unemployed in UK, 2013, Q4 (Office for National Statistics (2015b))
j	Job finding rate	0.207	Statistics on off-flow of JSA claimants UK, 2013, Q4 (Office for National Statistics (2015a)) 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
Repl. 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 (OECD (2014))
X	Number of observed vacancies	21.5	Half of the mean number of saved vacancies in our experiment
Extended model with On-The-Job search:			
$\tilde{j}$	Job-to-job transition rate	0.0082	Quarterly job-to-job transitions (UK, 2013-Q4) are 722,000 (Office for National Statistics (2013b)) while employment equals 30,288,000 (Office for National Statistics (2013a))
ρ	Probability of switching jobs conditional on an offer ($w = w^*$)	1	Set to 1.
$\tilde{\rho}$	Probability of switching jobs conditional on an offer ($w = y$)	0	Set to 0.

We use that the average number of vacancies the participants in our study saved was 9.9 per week and assume that about half of those are really of interest, thus setting $X = 9.9 \cdot 0.5 \cdot 4\frac{1}{3} = 21.5$.

6.3 Predictions

Now all elements necessary to compute r, w^*, $\lambda'(w^*)$ and the elasticity are known. Note that none of these values are targeted in our calibration; instead they are model predictions that arise after we calibrated the parameters listed in Table 9. We present the calibration results in Table 10. We find that the equilibrium wage is 0.805 which is slightly above the reservation wage.

In order to keep workers indifferent between the equilibrium wage and a deviation wage that is 20% higher, the queue length in this model has a wage elasticity of 18.7, far higher than the estimated

Table 10: Model predictions: No on-the-job-search

	Model parameter	Calibrated value	Experimental estimate
Equilibrium outcomes:			
Queue length	$\lambda(w^*)$	42.6	
Prob. suitable application	κ	0.065	
Equilibrium wage	w^*	0.805	
Reservation wage	r	0.802	
Unemployment benefits	b	0.49	
Employment at low wage	η_w	40.4 million	
Employment at high wage	η_y	4.5 million	
Key predictions:			
Wage elasticity	$\frac{w^*\lambda'(w^*)}{\lambda(w^*)}$	5.34	0.92
$P(A_h = 0 \mid A_l = 1)$	p_c	0	0.42
$P(A_l = 0 \mid A_h = 1)$	q_c	0.54	0.52

All parameter values are as described in Table 9.

elasticity in our experiment (0.92, presented in the third column of Table 3). In fact, given our assumptions on the number X of options that workers' see, the queue length at the deviant would not be given by indifference but rather all workers who observe the deviation would apply there. This still gives a model implied elasticity of 5.34, still an order of magnitude higher than our empirical counterpart.[49] Because the high wage is so attractive that all workers who see it apply there, we obtain a probability $p_c = 0$ that a worker who applies to a low wage job and sees the high wage job does not apply there. Obviously this is far from our empirical observation that a large fraction of workers only apply to the low wage.

This stark failure of this very basic model is not too surprising given the previous literature that we discussed in the introduction and review again in the next subsection. There we review possible reasons for this failure, and discuss which omissions from the model might bring it closer to the data. We then explore one particular prominent solution quantitatively.

6.4 Discussion

What is behind the extremely high elasticity of the queue length in the simple model? When a worker accepts a job, his wage is assumed to remain unchanged for the entire employment spell. In the data, employment spells last years in expectation, and the chance of finding another job next month conditional on remaining unemployed is high. So a 20% higher wage means a very high increase in expected net present value which makes it very attractive to apply, while the opportunity cost of extended unemployment if the application fails is relatively low. Since workers are assumed to fully believe the wage announcements, this explains the very high application elasticity in the simple model.

In fact, this feature should not be surprising given the preceding literature. As mentioned in the

[49] Indifference between the equilibrium wage and deviant wage requires a queue length at the higher wage of 201 (implying an elasticity of 18.7). The maximal queue length with imperfect observability is $Xu/v = 88.1$. With this maximal queue length the implied elasticity is $42.6 + El * 0.2 * 42.6 = 88.1$, or $El = 5.34$.

introduction, Hornstein et al. (2011) - hereafter HKV - show that a simple directed search model with reasonable parameters has a spectacularly high elasticity of the queue length. They use this to show that, starting from the average wage, the queue length would fall to zero even for very small decreases in the wage. Based on this insight they conclude that such models cannot support wage dispersion, as nobody would apply to wages that are even minor downward deviations from the average. Our finding is related, but in the direction of positive wage differences. We show that a higher wage would lead to a queue length vastly in excess of our empirical counterpart. HKV document a failure to sustain wage dispersion not only in directed search, but also in their main model which is based on random search. For their main model they explore one-by-one a large number of alternative specifications: European instead of U.S. labor market parameters, an additional "inactivity" state beyond employment and unemployment, endogenous search effort choice, psychological costs of unemployment, compensating differentials, stochastic wages, returns to experience, risk aversion with and without self-insurance, and on-the-job search. They argue that none of these resolves the problem, except: "Notable exceptions are some of the most recent models of on-the-job search" (HKV, abstract). They in particular single out models where the current employer bids against outside competitors that their worker meets on the job, as in Postel-Vinay and Robin (2002). In terms of theory it is an obvious fit to our baseline model, as our firms also bid against each other if several pursue the same worker. Since the problem in HKV is closely related to ours, and since they single out this type of on-the-job search as a remedy, we explore it quantitatively in our setting. This is presented in the next subsection, and turns out to reconcile the model closely with the empirical counterparts.

Before exploring this quantitatively, we note that this obviously is not the only omission from the model that has the potential of reconciling it with our empirical findings. In our base model the starting wage of a job characterizes its wage throughout the job spell, and therefore its net present value. If jobs are characterized by wage-tenure contracts with back-loaded wages, then the starting wage is not a good indicator of the net present value. In particular, if wage progression is faster when the initial wage is lower, the net present value of two jobs is closer together than their starting wages indicate. So initial wage differences are not so important if the low wage job catches up over time and, therefore, the queue length does not react that much. This arises for example in sophisticated on-the-job search models such as Burdett and Coles (2003) or Shi (2009), where bidding against outside offers is assumed impossible and firms use back-loaded wage-tenure contracts to lower attrition. HKV argues that quantitatively such models are similar to models with bidding in the spirit of Postel-Vinay and Robin (2002), since in either version wages increase over time and more so for the job that initially pays less. We expect that this indeed offers an alternative route to reconcile the model with our empirical findings.

Our model follows most of the directed search literature by assuming a tight connection between the posted wage and the starting wage once the worker is hired. In reality this connection might be more tenuous. Workers might assume some wage bargaining to happen even once they applied to the firm, breaking this tight link. One way in which this has been explored theoretically is to assume that wages are just cheap talk that proceeds a subsequent wage negotiation. Menzio (2007) shows that firms that differ in productivity can still use such cheap talk to send some information about their productivity to workers, which will allow better firms to attract more workers albeit at the cost of paying higher wages at the bargaining phase. This model has a much more tenuous link between cheap talk wages offered ex ante, and wages that are paid ex post. While we are not aware of a quantitative attempt to discipline this, it has the potential to lower the queue length elasticity to the level of our empirical findings.

Alternatively, one could interpret the wage postings as minimum wages that firms commit to pay, but that leaves them the freedom to improve their offer later when they actually meet the worker. Since wage postings are less tightly linked to actual wages, this might lower the elasticity of the queue length with respect to posted wages. Shi (2019) formulates such a setup where workers target a particular minimum wage posted by firms, but might meet multiple firms with identical minimum wage.[50] Because a worker might have alternative offers, it turns out that firms are willing to raise their wage above the posted minimum.[51] Shi (2019) does not provide a quantitative assessment, but the mechanism is related to features that are already present in our calibration above: also in our model firms bid up the wage above the posted wage if the worker has multiple offers. A difference is that firms in our model observe whether a worker has met another firm, while they do not observe this in Shi (2019). Still, this similarity also breaks the tight link between wage announcement and final wage payment in our baseline model. This turn out to be insufficient to generate quantitatively plausible magnitudes in the previous section, which might indicate a limited potential for this route overall.

Our model does not account for the presence of a large share of vacancies that do not provide wage information. Michelacci and Suarez (2006) formulate a reason why some firms in the market might post wages and others post none, as is the case in our empirical setting even though our experiment focuses only on the part that actually posts wages. They present a model in which workers differ in productivity, but wage offers cannot condition on this. Workers' productivity can only be verified on the job, and enters the bargaining outcome. Bargaining has the advantage of adjusting the wage to some extent to workers' productivity, while wage posting has the advantage of attracting workers to ones own firm. The setting can give rise to both wage posting and bargaining. While it can reconcile that only some firms post wages, this type of model is unlikely to resolve the excessive queue length elasticity: for those firms that do offer wages this is a commitment, and we expect the queue length to react as strongly as in our baseline model. We are not aware of a quantitative exploration of this theory, though.

Finally, we assume that job seekers pay full attention to the wage differences between vacancies. We abstracts from mistakes in their behavior as well as from inattention. Either of these could introduce an element of random search into this environment, implying that search behavior is only partially directed. This would obviously dampen the reaction of the queue length with regards to the wage. Lentz and Moen (2017) outline a model in which agents make mistakes that decrease in the utility loss relative to the optimal action, allowing them to span the space between random and directed search. While this has the potential to dampen the reaction of the queue length, this has not yet been applied

[50] We thank an anonymous referee for bringing this work to our attention.

[51] The model of Shi (2019) has two wages: the posted minimum and the subsequent updated wage offer. It is difficult to map either to the wages in our experiment. If the experimental wage represents the minimum wage in their theory, their assumption that each worker can only pursue one such minimum wage is hard to interpret given that in the experiment many workers are interested in both wages. Extending their theory to allow the pursuit of multiple of these minimum wages is non-trivial, and no theory work exists to build upon. If the experimental wage posting corresponds to the updated wage offer in their theory, then workers in their theory are sure to get the job and always accepts the best offer, while in our setting workers are clearly unsure whether they will be hired.

in a quantitative way. Wu (2020) provides an extremely elegant model of rational inattention that attaches a cost to deviations from random search. Again, this can span the entire space from directed to random search. In our empirical setting we go to great length to ensure that agents have had both of our research vacancies displayed to them, which we can ensure because we run the search engine. But since we did not use sophisticated tools such as eye tracking software, we cannot ensure that they actually truly see both vacancies. Even if they did, they might not fully digest all information. While Wu (2020) does not explore whether the amount of inattention that is necessary to generate a reasonable elasticity of the queue length has plausible empirical properties, he does build on our work to illustrate how the equilibrium price would be affected if inattention were high enough to generate our empirical magnitude.

To our knowledge, none of the papers in the literature has considered whether those who apply to the low wage will apply also to a high wage. That is due to the fact that prior to our work we lacked empirical evidence on this, and due to the common model simplification that considers a single job application per worker.

In the following we will augment our basic multiple-application model with the most successful extension from HKV: we introduce on-the-job search where wages are bid up when two firms pursue the same worker.

6.5 On-the-job search

On-the-job search (OJS) can be introduced into the model in a fairly straightforward way. We relegate the model to the appendix and briefly describe the additions here. Workers employed at w^* are able to search for jobs and we assume that firms observe whether applicants are already employed elsewhere. In case an employed worker secures a job offer, his wage is bid up to productivity by his current and new employer. Firms can only make one job offer per period, and always prefer to offer the job to an unemployed worker because an employed worker would trigger a bidding war that leaves no surplus to the firm. From the perspective of an unemployed worker, this means that s(he) competes only with other unemployed workers for any given job. Workers employed at wage $w < y$ who get an outside offer move to the new firm with probability ρ. Workers employed at $w = y$ are equally likely to attract job offers, but since they have no further gains, we assume that they move to a new job (conditional on the offer) with probability $\tilde{\rho} < \rho$. In the baseline calibration we set $\rho = 1$ and $\tilde{\rho} = 0$.

Calibration of the model with OJS is largely similar to the model without OJS. We set the job-to-job transition rate ($\tilde{j}$) equal to 0.19%, which is the job-to-job transition rate in the UK in the fourth quarter of 2013 (ONS statistics, see lower panel of Table 9). Using the value of ($\tilde{j}$), we can back out the probability of receiving an offer for an employed worker ($\tilde{m}$) using steady state conditions for the number of employed workers at either $w = w^*$ and $w = y$ (see appendix A.5 for details).

The results from calibrating the model with OJS are presented in the second column of Table 11. As expected, we find that adding OJS leads to a substantially lower wage elasticity of 0.57. While somewhat smaller than our experimental finding (in the range of 0.7 to 0.92), it clearly has the same order of magnitude. Thus, a directed search model can produce an estimate of job seekers' responsiveness to posted wages that is close to our experimental finding. Given this reasonable elasticity, we also find sensible values for the probability p_c (q_c) of applying to a low (high) wage vacancy but not to

Table 11: Model predictions: With on-the-job-search

	Model parameter	Calibrated value	Experimental estimate
Equilibrium outcomes:			
Queue length	$\lambda(w^*)$	42.6	
Prob. suitable application	κ	0.065	
Equilibrium wage	w^*	0.100	
Reservation wage	r	0.087	
Unemployment benefits	b	0.525	
Employment at low wage	η_w	6.2 million	
Employment at high wage	η_y	38.7 million	
Key predictions:			
Wage elasticity	$\frac{w^*\lambda'(w^*)}{\lambda(w^*)}$	0.57	0.92
$P(A_h = 0 \mid A_l = 1)$	p_c	0.49	0.42
$P(A_l = 0 \mid A_h = 1)$	q_c	0.54	0.52

All parameter values are as described in Table 9.

the high (low) wage vacancy. We find that $p_c = 0.48$ and $q_c = 0.54$. The corresponding probabilities obtained from the experiment are 0.42 (see Table 5) and 0.52 (see Table 12 in the Appendix). The model again produces values very close to the empirical findings. It nicely captures the fact that the difference between applying-low-but-not-high and the reverse is quite small. We should note that the predicted probabilities are highly dependent on the value of X, the number of vacancies that a job seeker observes (although the difference between p_c and q_c is fairly stable). In line with expectations, if the set of observed vacancies is larger, the probability of not applying to one in particular increases (that is, both p_c and q_c increase). Since we have no direct evidence on the number of vacancies that are of interest, we show how the choice of X affects predictions in appendix A.6.

These results indicate that a simple parametrized directed search model generates a wage elasticity that is 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. Using reasonable calibration values, we are able to rationalize this nearly perfectly, even though at first sight our empirical finding might have been puzzling. We kept the setup deliberately simple, abstracting from effects of starting wages on future wages or wage improvements in the current job beyond those induced through bidding among multiple potential employers. Such features could be introduced in future work. The success of our simple model gives hope that even more complex models in this model class might be 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.

7 Conclusion

In this study we present results on how the wage announcement in vacancies affects behavior of potential applicants. By posting pairs of vacancies with randomly assigned wages we provide evidence that, holding all vacancy characteristics constant, a higher wage attracts more applicants. This contrasts with the common finding that in observational data higher wage jobs receive fewer applications, even when controlling for occupation. This finding shows that wage postings affect job seeker behavior, suggesting that job search is directed rather than random.

We investigate three extensions that relate to predictions from directed search theory. First, a non-trivial share of applicants only shows interest in a low wage vacancy even though a high-wage vacancy, which is otherwise identical, exists. This is consistent with job seekers sending only a limited amount of job applications strategically, and trading-off the wage with the probability of getting hired. Second, a complementary survey shows that higher wages are perceived to reflect more competition for the job (both in quality and quantity of applicants) while higher wages also reflect better non-working conditions. The fact that wage differences affect beliefs on competition fits with the key mechanism of directed search. Third, we calibrate a simple homogeneous-agent directed search model and study its ability to replicate these findings. The simplest model where the announced wage lasts the entire employment spell misses the empirical observations by an order of magnitude. When coupled with on-the-job search it replicates the application-elasticity well, and allowing for multiple applications per individual captures the fact that many individuals apply only to the low wage. Jointly, the findings indicate that a directed search approach seems promising to study the reaction of job seekers to changes in posted wage offers.

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. A particularly fruitful application could investigate those vacancies that do not post wage offers. Our study is silent on those, as it focuses on variations of wages for those vacancies that already use wage offers. For those that do not, it might be interesting to explore what happens if one randomly introduces wage offers. Alternatively, one could randomly vary those words or phrases that are empirically associated with higher wages in the cross-section, to see if they might serve as a proxy for wage offers.

References

Abebe, G., Caria, A. S., and Ortiz-Ospina, E. (2021). The selection of talent: Experimental and structural evidence from Ethiopia. *American Economic Review*, 111(6):1757–1806.

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. (2020). Directed search with phantom vacancies. IZA Discussion Paper No. 7403.

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

Anbarci, N. and Feltovich, N. (2013). How sensitive are bargaining outcomes to changes in disagreement payoffs? *Experimental economics*, 16(4):560–596.

Ashraf, N., Bandiera, O., Davenport, E., and Lee, S. S. (2020). Losing prosociality in the quest for talent? Sorting, selection, and productivity in the delivery of public services. *American Economic Review*, 110(5):1355–94.

Banfi, S. and Villena-Roldan, B. (2019). Do high-wage jobs attract more applicants? Directed search evidence from the online labor market. *Journal of Labor Economics*, 37(3):715–746.

Belot, M., Kircher, P., and Muller, P. (2019). Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice. *The Review of Economic Studies*, 86(4):1411–1447.

Belot, M., Kircher, P., and Muller, P. (2021). Data and code for: How wage announcements affect job search - a field experiment. Technical report, American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. http://doi.org/10.3886/E152721V1.

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.

Boal, W. M. and Ransom, M. R. (1997). Monopsony in the labor market. *Journal of economic literature*, 35(1):86–112.

Braun, C., Engelhardt, B., Griffy, B., and Rupert, P. (2016). Do workers direct their search? mimeo.

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 Coles, M. (2003). Equilibrium wage-tenure contracts. *Econometrica*, 71(5):1377–1404.

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

Burdett, K., Shi, S., and Wright, R. (2001). Pricing and matching with frictions. *Journal of Political Economy*, 109(5):1060–1085.

Cason, T. N. and Noussair, C. (2007). A market with frictions in the matching process: an experimental study. *International Economic Review*, 48(2):665–691.

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.

Cooley, T., Marimon, R., and Quadrini, V. (2020). Commitment in organizations and the competition for talent. *Review of Economic Studies*, 87(5):2165–2204.

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.

Deserranno, E. (2019). Financial incentives as signals: Experimental evidence from the recruitment of village promoters in uganda. *American Economic Journal: Applied Economics*, 11(1):277–317.

Dube, A., Jacobs, J., Naidu, S., and Suri, S. (2020). Monopsony in online labor markets. *American Economic Review: Insights*, 2(1):33–46.

Engelhardt, B. and Rupert, P. (2017). Competitive versus random search with bargaining: An empirical comparison. *Labour Economics*, 48:183–197.

Faberman, R. J. and Menzio, G. (2018). Evidence on the relationship between recruiting and the starting wage. *Labour Economics*, 50:67–79.

Falch, T. (2010). The elasticity of labor supply at the establishment level. *Journal of Labor Economics*, 28(2):237–266.

Farber, H. S., Silverman, D., and von Wachter, T. (2016). Determinants of callbacks to job applications: An audit study. *American Economic Review*, 106(5):314–318.

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.

Gurio, F. and Rudanko, L. (2014). Customer capital. *Review of Economic Studies*, 81(3):1102–1136.

Helland, L., Moen, E. R., and Preugschat, E. (2017). Information and coordination frictions in experimental posted offer markets. *Journal of Economic Theory*, 167:53–74.

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.

Kloosterman, A. (2016). Directed search with heterogeneous firms: an experimental study. *Experimental Economics*, 19(1):51–66.

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.

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.

Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press, London.

Manning, A. (2011). Imperfect competition in the labor market. volume 4, Part B of *Handbook of Labor Economics*, pages 973 – 1041. Elsevier.

Manning, A. (2021). Monopsony in labor markets: A review. *ILR Review*, 74(1):3–26.

Marinescu, I. (2017). The general equilibrium impacts of unemployment insurance: Evidence from a large online job board. *Journal of Public Economics*, 150:14–29.

Marinescu, I. and Rathelot, R. (2018). Mismatch unemployment and the geography of job search. *American Economic Journal: Macroeconomics*, 10(3):42–70.

Marinescu, I. and Wolthoff, R. (2020). Opening the black box of the matching function: The power of words. *Journal of Labor Economics*, 38(2):535–568.

Mas, A. and Pallais, A. (2017). Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–59.

McCall, J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 84:113–126.

Menzio, G. (2007). A theory of partially directed search. *Journal of Political Economy*, 115(5):748–769.

Menzio, G. and Shi, S. (2011). Efficient search on the job and the business cycle. *Journal of Political Economy*, 119:468–510.

Michelacci, C. and Suarez, J. (2006). Incomplete wage posting. *Journal of Political Economy*, 114(6):1098–1123.

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.

OECD (2014). OECD statistics: Net replacement rate in unemployment. https://stats.oecd.org/Index.aspx?DataSetCode=NRR. Accessed: 2016-04-26.

Office for National Statistics (2011-2015b). Nomis Official Labour Market Statistics: Model-based estimates of unemployment (City of Edinburgh and UK). https://www.nomisweb.co.uk/datasets/umb. Accessed: 2017-09-14.

Office for National Statistics (2012-2015a). Nomis Official Labour Market Statistics: Claimant Count (City of Edinburgh and UK). https://www.nomisweb.co.uk/sources/cc. Accessed: 2015-03-27.

Office for National Statistics (2013a). UK Employment, seasonally adjusted. https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/timeseries/mgrz/lms. Accessed: 2016-04-28.

Office for National Statistics (2013b). UK Job-to-Job moves. https://www.ons.gov.uk/ employmentandlabourmarket/peopleinwork/employmentandemployeetypes/adhocs/ 004892totaljobtojobmovesjulytoseptember2015. Accessed: 2016-04-28.

Office for National Statistics (2013c). UK Vacancies. https://www.ons.gov.uk/ employmentandlabourmarket/peopleinwork/employmentandemployeetypes/timeseries/ap2y/unem. Accessed: 2016-04-28.

Peters, M. (1991). Ex ante price offers in matching games: Non-steady states. *Econometrica*, 59:1425–1454.

Peters, M. (2010). Noncontractible heterogeneity in directed search. *Econometrica*, 78:1173–1200.

Postel-Vinay, F. and Robin, J.-M. (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica*, 70(6):2295–2350.

Rosen, S. (1986). The theory of equalizing differences. In Ashenfelter, O. and Layard, R., editors, *Handbook of Labor Economics*, volume I, pages 641–692. Elsevier Science Publishers.

Schaal, E. (2017). Uncertainty and unemployment. *Econometrica*, 85(6):1675–1721.

Shi, S. (2002). A Directed Search Model of Inequality with Heterogeneous Skills and Skill-Biased Technology. *The Review of Economic Studies*, 69(2):467–491.

Shi, S. (2009). Directed search for wage tenure contracts. *Econometrica*, 77(2):561–584.

Shi, S. (2019). Sequentially mixed search and equilibrium price dispersion. SSRN working paper no 3262935.

Shimer, R. (2005). The assignment of workers to jobs in an economy with coordination frictions. *Journal of Political Economy*, 113(5):996–1025.

Staiger, D. O., Spetz, J., and Phibbs, C. S. (2010). Is there monopsony in the labor market? Evidence from a natural experiment. *Journal of Labor Economics*, 28(2):211–236.

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. (2019). Directed search: A guided tour. *Journal of Economic Literature, forthcoming.*

Wu, L. (2020). Partially directed search in the labor market. Mimeo.

A Appendix

A.1 Institutional Setting

In Figure 4 we present aggregate labor market statistics (Office for National Statistics (2015b), Office for National Statistics (2015a)). 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) Unemployment rate

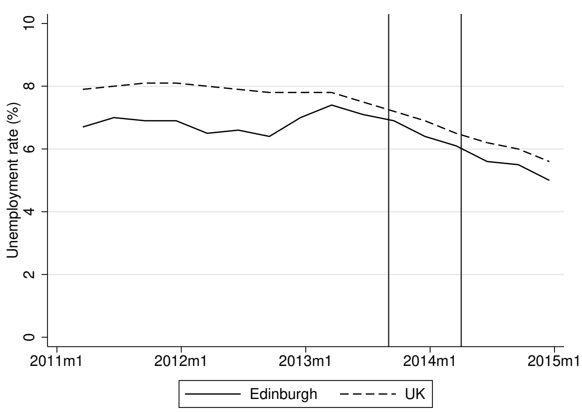

(b) JSA claimants

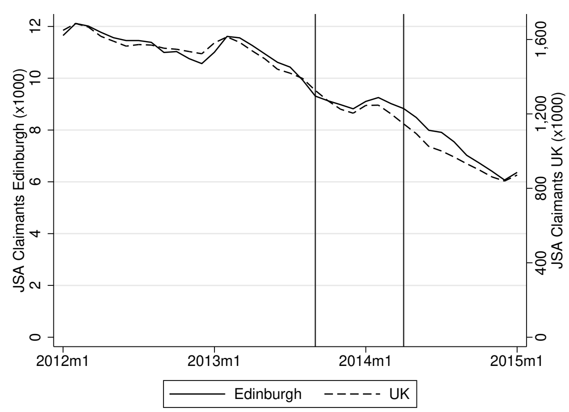

A.2 Job Search Interface

Figure 5: Screenshot of the standard job search interface

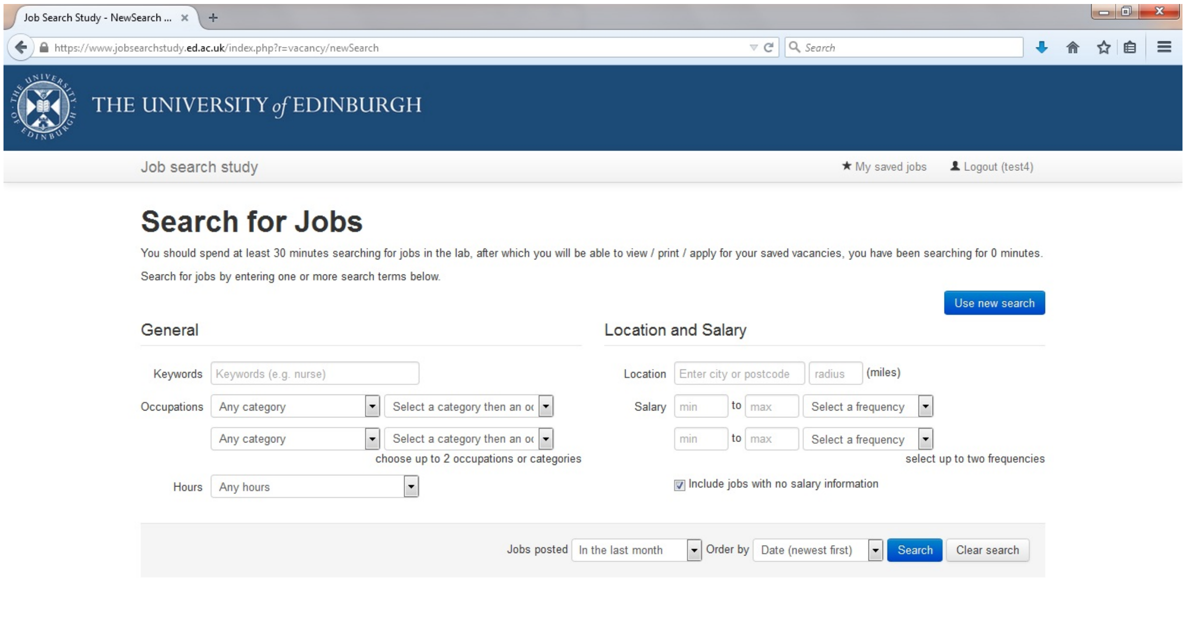

A.3 Detailed robustness checks for the reservation wage property results

Table 12: Probability of not saving low-wage, conditional on saving high-wage

	Saves			Views		
	(1)	(2)	(3)	(4)	(5)	(6)
Low appears first			-0.20**			-0.17***
			(0.062)			(0.047)
Constant	0.54***	0.52***	0.60***	0.50***	0.47***	0.54***
	(0.034)	(0.035)	(0.043)	(0.024)	(0.025)	(0.028)
Sample	All	List both	List both	All	List both	List both
N	337	318	318	623	590	590

* p < 0.10, ** p < 0.05, *** p < 0.01. Standard errors in parentheses.

Table 13: Probability of saving low but not high: robustness checks (1)

	Exclude posted twice		Exclude based on exit survey		Exclude based on distance	
	(1)	(2)	(3)	(4)	(5)	(6)
High appears first		-0.26***		-0.39***		-0.34***
		(0.056)		(0.074)		(0.070)
Constant	0.39***	0.49***	0.36***	0.48***	0.48***	0.59***
	(0.030)	(0.040)	(0.047)	(0.059)	(0.039)	(0.041)
Sample	List both	List both	List both	List both	List both	List both
N	252	252	116	116	135	135

* p < 0.10, ** p < 0.05, *** p < 0.01. Clustered standard errors (by pair) in parentheses.

Table 14: Probability of viewing low but not high: robustness checks (1)

	Exclude posted twice		Exclude based on exit survey		Exclude based on distance	
	(1)	(2)	(3)	(4)	(5)	(6)
High appears first		-0.29***		-0.30***		-0.23***
		(0.047)		(0.062)		(0.060)
Constant	0.41***	0.52***	0.38***	0.48***	0.39***	0.47***
	(0.029)	(0.034)	(0.038)	(0.051)	(0.036)	(0.045)
Sample	List both	List both	List both	List both	List both	List both
N	463	463	218	218	256	256

* p < 0.10, ** p < 0.05, *** p < 0.01. Clustered standard errors (by pair) in parentheses.

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, and we emphasize that our main results for viewing have been very similar to those for saving.

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 Tables 13 and 14, columns (1) and (2). We find that removing the 74 observations has very little effect on the probabilities: the non-saving probabilities decrease slightly, while the non-viewing probabilities increase slightly. 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 Tables 13 and 14, columns (3) and (4). The probabilities of not saving or viewing the high wage vacancy conditional on doing so for the low wage one 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 0.09). 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. 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 0.16. 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. A 95 % confidence interval includes values between 0.03 and 0.23 and thus the number is not significantly different from our previous estimates, while it is significantly different from zero.[52]

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

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

Table 15: Probability of saving low but not high: robustness checks (2)

	Exclude if lists many		Exclude large salary difs	
	(1)	(2)	(3)	(4)
High appears first		-0.26***		-0.25**
		(0.072)		(0.078)
Constant	0.30***	0.41***	0.41***	0.48***
	(0.039)	(0.053)	(0.038)	(0.046)
Sample	List both	List both	List both	List both
N	110	110	132	132

* p < 0.10, ** p < 0.05, *** p < 0.01. Clustered standard errors (by pair) in parentheses.

Table 16: Probability of viewing low but not high: robustness checks (2)

	Exclude if lists many		Exclude large salary difs	
	(1)	(2)	(3)	(4)
High appears first		-0.29***		-0.21**
		(0.059)		(0.061)
Constant	0.31***	0.43***	0.41***	0.48***
	(0.032)	(0.046)	(0.034)	(0.045)
Sample	List both	List both	List both	List both
N	194	194	260	260

* p < 0.10, ** p < 0.05, *** p < 0.01. Clustered standard errors (by pair) in parentheses.

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.[53] 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 13 and 14, columns (5) and (6). 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

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

vacancies per week into half (around 250 per week). Individuals that list relatively few vacancies have 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.[54] The results are presented in Table 15 and 16, columns (1) and (2). We find that the probability of not saving the high wage vacancy, while saving the low wage vacancy is somewhat lower for this group. When both vacancies have been listed this probability is 0.31, while it was 0.42 in the complete sample. This is a natural result that also follows from the directed search model that we discuss and calibrate in section 6: the smaller the set of potentially interesting vacancies, the larger the probability of saving each particular vacancy. A similar difference is found when considering reverse probabilities: the probability of not saving the low one when saving the high one is 0.43 for this sample while it is 0.52 for the complete sample (not reported in the tables). Moreover, if we restrict the sample further, this probability does not appear to decrease further.[55] In conclusion, it appears that some individuals prefer the lower wage vacancy, as even within this sample a substantial share of individuals does not save the high wage vacancy while saving the low wage vacancy.

Within pair wage difference The within pair wage differences that we imposed vary from 20 to 40 %. One concern with the larger wage difference is that a very high or very low wage might look suspicious to job seekers, causing them to ignore such vacancies. While our check of the real vacancy wage support (Figure 2) is reassuring, we can simply perform a robustness check by considering only pairs with a smaal (20 %) wage difference. We find a very similar probability for not saving the high wage vacancy while saving the low wage vacancy (0.41 vs 0.42 in the baseline). Also the viewing probabilities are very similar, suggesting that the wage range does not drive our saving and viewing results.

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 artificial vacancies decreased slightly over the 12 weeks, such that absolute numbers are uninformative. Alternatively we compute the ratio of savings per viewing *for artificial vacancies*. If participants get better at identifying artificial vacancies (once they view one) this ratio would be expected to decrease over time. We find that it is 0.53 in weeks 1-4, 0.51 in weeks 5-8 and 0.58 in weeks 9-12. Two-sided t-tests for equality show that none of these ratios differ significantly from each other. In addition, 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 any differences in the saving rate to differences in the pool of artificial vacancies and differences in the pool of participants (due to attrition).

[54] Note that listing fewer vacancies is not correlated with a shorter time period spent in the experiment sessions on average, and fewer listings thus really suggest more time spent on each listing.

[55] When restricting the sample to individuals that list less than 200 vacancies per week, we find the corresponding probability is 0.36 (as compared to 0.31). Such a selection includes only 28 observations and thus the results become rather imprecise.

A.4 Equilibrium Derivation

In this section we derive the model equilibrium from section 6 and the reaction to deviations, for the case with on-the-job search. Note that on-the-job search can simply be shut down by setting $\tilde{m} = 0$.

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

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

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:

$$\begin{aligned} \Pi(w) &= y - w + \beta\delta(1 - \tilde{m}^*)\Pi(w) \\ &\Leftrightarrow \quad \Pi(w) = \frac{y - w}{1 - \beta\delta(1 - \tilde{m}^*)}, \end{aligned} \tag{11}$$

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.[56] AGV show in a model without on-the-job search that the equilibrium offered wage w^* falls to the worker's reservation value r at which the worker is exactly indifferent between accepting the job and not accepting (they apply nevertheless because they hope to get two offers, in which case they can bid up the wage). 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:

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

[56] This 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.

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

$$\begin{aligned} U^* &= \frac{b + Nm^*[1-m^*]^{N-1}\beta V(w^*)}{1-[1-m^*]^N\beta} \\ &\quad + \frac{[1-[1-m^*]^{N-1}[1+(N-1)m^*]]\beta V(y)}{1-[1-m^*]^N\beta}. \end{aligned} \tag{12}$$

where $V(.)$ 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

$$V(y) = \frac{y + \beta(1-\delta)U^*}{1-\delta\beta}. \tag{13}$$

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:

$$\begin{aligned} V(w) &= w + \beta\delta\tilde{m}^*V(y) + \beta\delta(1-\tilde{m}^*)V(w) + \beta(1-\delta)U^* \\ &\Leftrightarrow \quad V(w) = \frac{w + \beta\delta\tilde{m}^*V(y) + \beta(1-\delta)U^*}{1-\delta\beta(1-\tilde{m}^*)}. \end{aligned} \tag{14}$$

Jointly, equations (12), (13) and (14) 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 (11)). The first-order condition is

$$\begin{aligned} & n'(\lambda(w^*))\lambda'(w^*)[1-m(\lambda(w^*))]^{N-1}\beta\frac{y-w^*}{1-\beta\delta(1-\tilde{m}^*)} \\ & \quad -n(\lambda(w))[1-m(\lambda(w^*))]^{N-1}\beta\frac{1}{1-\beta\delta(1-\tilde{m}^*)} \quad \leq \quad 0 \end{aligned} \tag{15}$$

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

$$
\begin{aligned}
U(w^d) \quad = \quad & b+m(\lambda(w^d))[1-m^*]^{N-1}\beta V(w^d) \\
& +(N-1)m^*[1-m(\lambda(w^d))][1-m^*]^{N-2}\beta V(w^*) \\
& +[1-m(\lambda(w^d))][1-m^*]^{N-1}\beta U^* \\
& +\begin{bmatrix} 1-m(\lambda(w^d))[1-m^*]^{N-1} \\ -(N-1)m^*[1-m(\lambda(w^d))][1-m^*]^{N-2} \\ -[1-m(\lambda(w^d))][1-m^*]^{N-1} \end{bmatrix}\beta V(y)
\end{aligned} \tag{16}
$$

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

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

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:

$$
\begin{aligned}
\lambda(w^d) &= p\frac{X}{v}u \\
p &= \frac{\lambda(w^d)v}{Xu}
\end{aligned} \tag{17}
$$

The value p should be such that workers are indifferent between applying to the deviant vacancy and equilibrium wage vacancies. Using equation 6 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:

$$q = \frac{N-p}{X-1} \tag{18}$$

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 $p_c \equiv \Pr(A_h = 0|A_l = 1) = \frac{N(1-p)}{N-p}$.[57] 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.5 Steady state conditions (with job-to-job transitions)

We calibrate the on-the-job matching rate $\tilde{m}$ from the empirical job-to-job transition rate ($\tilde{j}$) using steady state conditions for the number of workers employed at the low wage n_w and the high wage n_y:

$$\Delta n_w = uNm^*(1-m^*)^{N-1} - n_w(1-\delta+\tilde{m}\delta) \tag{20}$$

In steady state $\Delta n_w = 0$ such that we have

$$n_w = \frac{uNm^*(1-m^*)^{N-1}}{(1-\delta(1-\tilde{m})} \tag{21}$$

Similarly, the number of employed at a wage equal to productivity is

$$n_y = \frac{n_w\delta\tilde{m} + u\Big[1 - Nm^*(1-m^*)^{N-1} - (1-m^*)^N\Big]}{1-\delta} \tag{22}$$

Finally, the aggregate empirical job-to-job transition rate $\tilde{j}$ equals the total number of job switches divided by the number of employed workers:

$$\tilde{j} = \frac{\delta\tilde{m}(n_w\rho + n_y\tilde{\rho})}{n_w + n_y} \tag{23}$$

Where, conditional on an offer, a worker at the low wage moves to the new job with probability ρ, while for high-wage workers this probability is $\tilde{\rho}$. Substituting n_w and n_y using (21) and (22) and solving for $\tilde{m}$ gives the expression for the transition rate for employed workers.

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

$$\Pr(A_h = 0|A_l = 1) = \frac{\Pr(A_h = 0, A_l = 1)}{\Pr(A_l = 1)} = \frac{\Pr(A_l = 1|A_h = 0)\Pr(A_h = 0)}{\Pr(A_l = 1)} \tag{19}$$

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 (19) and using equation (18) to replace q we find $\Pr(A_h = 0|A_l = 1) = \frac{N(1-p)}{N-p}$.

Table 17: Sensitivity of model predictions to choice of X (number of observed vacancies)

Weekly # of observed vacancies	Empirical source	Monthly observed vacancies (X)	$P(A_h = 0 \mid A_l = 1)$	$P(A_l = 0 \mid A_h = 1)$
24	# of viewed	104	0.90	0.91
9.9	# of saved	42.9	0.75	0.78
8.3	# of saved (sample that lists few)	36	0.70	0.73
4.45	50% of the # of saved	21.5	0.49	0.54
4.15	50% of the # of saved (sample that lists few)	18	0.38	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 Number of observed vacancies in the model calibration

In the calibration of the directed search model (section 6.3), 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 15, 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 17 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.

A.7 Additional results

Table 18: Employer survey: number of observations

Surveyed employers (all vacancies in Scotland)	13,160
Responses	1014
Linked to vacancy data (only Edinburgh vacancies)	368
Has posted wage	177

This survey was sent out to all vacancies that included an email address and were posted in Scotland during the first wave of the experiment.

Table 19: Starting wage relative to posted wage

Point wage posted (66 vacancies)	
Below posted wage	9 (14%)
Equal to posted wage	37 (56%)
Above posted wage	20 (30%)
Wage range posted (111 vacancies)	
Below posted range	18 (16%)
Within posted range	77 (69%)
Above posted range	16 (14%)

Table 20: Effect of wage on number of views: real vacancies

			Poisson regression				Log-log regression
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log(Salary)	-0.65***	-0.61***	-0.50***	-0.50***	-0.49***	-0.16***	-0.45***
	(0.0082)	(0.011)	(0.015)	(0.015)	(0.041)	(0.018)	(0.043)
Temporary contract			-0.25***	-0.25***	-0.23**	-0.21***	-0.40***
			(0.021)	(0.021)	(0.080)	(0.020)	(0.047)
Part time			1.00***	1.00***	0.98***	0.41***	0.73***
			(0.060)	(0.060)	(0.22)	(0.035)	(0.065)
No company name			-0.44***	-0.49***	-0.46***	-0.10***	-0.43***
			(0.016)	(0.015)	(0.049)	(0.026)	(0.046)
No contacts in ad				0.46***	0.50***	0.33***	0.27***
				(0.047)	(0.16)	(0.033)	(0.053)
Constant	45546.2***						2.72***
	(10497.1)						(0.45)
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	7165	7010	6936	6936	6936	4962	7091

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. 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(views+0.1)

Table 21: Effect of wage on number of saves: real vacancies that were posted on the days on which we posted artificial vacancies (plus Mondays)

	Poisson regression						Log-log regression
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log(Salary)	-0.73***	-0.73***	-0.64***	-0.64***	-0.63***	-0.24*	-0.54***
	(0.014)	(0.041)	(0.059)	(0.057)	(0.062)	(0.12)	(0.10)
Temporary contract			-0.17	-0.16	-0.18	-0.33**	-0.32**
			(0.14)	(0.14)	(0.14)	(0.12)	(0.13)
Part time			0.74***	0.73***	0.72***	0.37*	0.68***
			(0.36)	(0.35)	(0.36)	(0.25)	(0.18)
No company name			-0.66***	-0.68***	-0.67***	-0.018	-0.59***
			(0.057)	(0.058)	(0.061)	(0.17)	(0.12)
No contacts in ad				0.31	0.30	0.37**	-0.037
				(0.24)	(0.23)	(0.20)	(0.14)
Constant	360686.3***						4.54***
	(180993.7)						(1.05)
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	2059	1873	1860	1860	1860	1281	2046

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

Figure 6: Salary differences in terms of standard deviations across 3-digit occupations (only real vacancies without company name)

(a) 20% wage difference

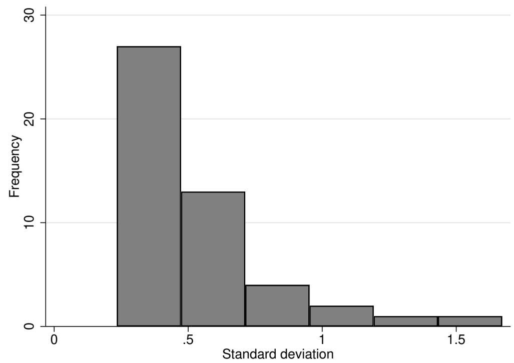

(b) 40% wage difference

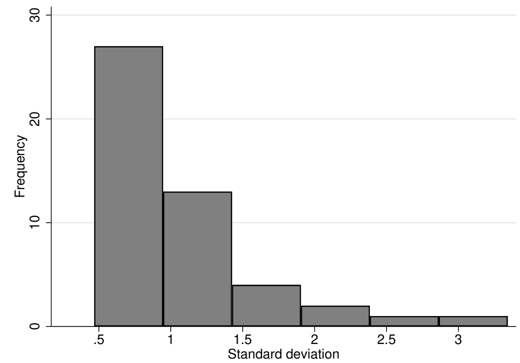

Figure 7: Number of times a vacancy was saved and viewed: Edinburgh pairs

(a) Saved

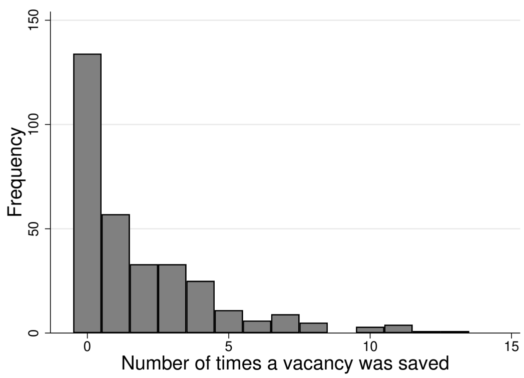

(b) Viewed

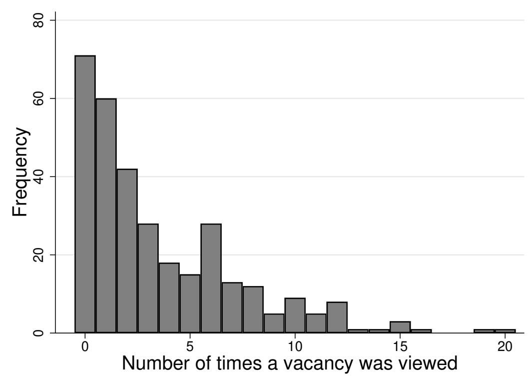

Figure 8: Number of times a vacancy was saved and viewed: Glasgow pairs

(a) Saved

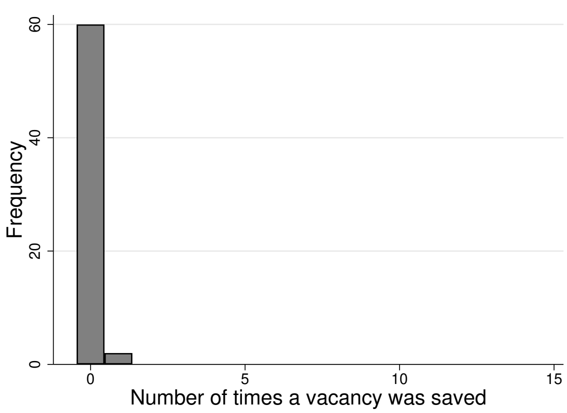

(b) Viewed

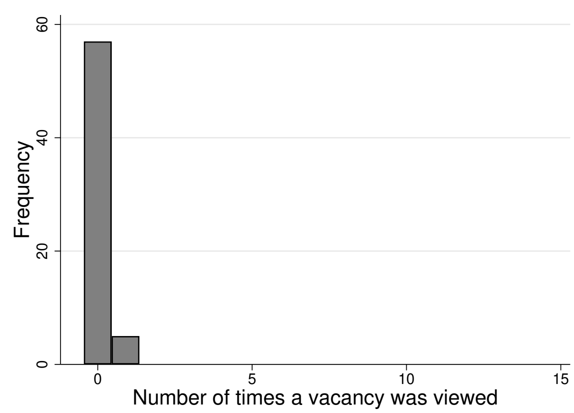

Figure 9: Cumulative distribution of the within pair difference between the higher and lower wage vacancies in the pairs

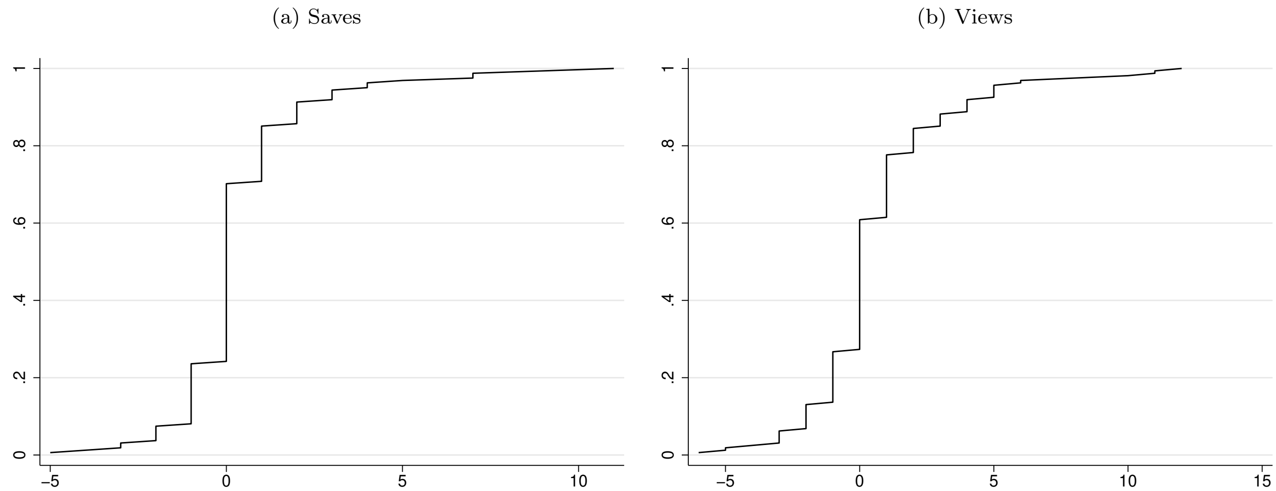

Table 22: Effect of wage difference on number of saves/views: Edinburgh and Glasgow pairs

	Saves			Views		
	(1)	(2)	(3)	(4)	(5)	(6)
Salary difference (in %)	0.73**	0.72**	0.96***	0.72***	0.73***	0.87***
	(0.45)	(0.46)	(0.44)	(0.36)	(0.36)	(0.30)
Appearing first			0.58***			0.50***
			(0.13)			(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).

Figure 10: Cronbach alpha's for the 16 vacancy sets

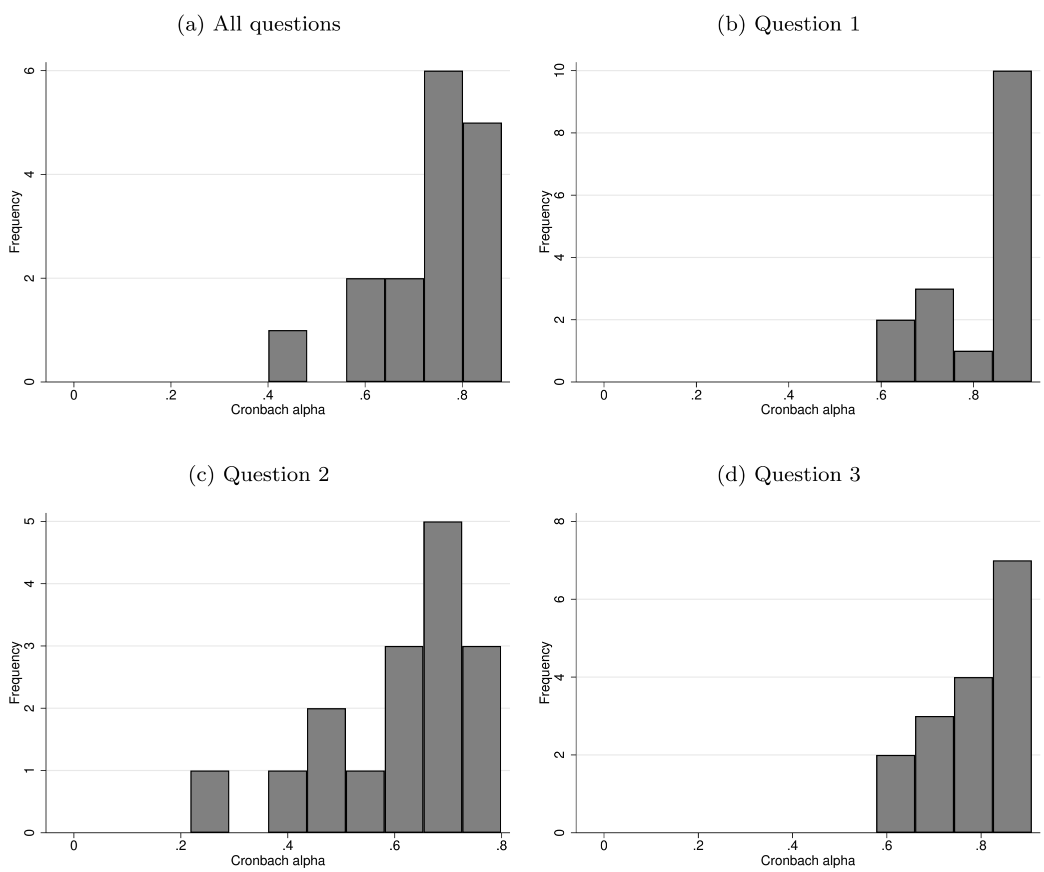

(a) All questions

(b) Question 1

(c) Question 2

(d) Question 3

Figure 11: Cronbach alpha's for the 161 vacancy pairs

(a) All questions

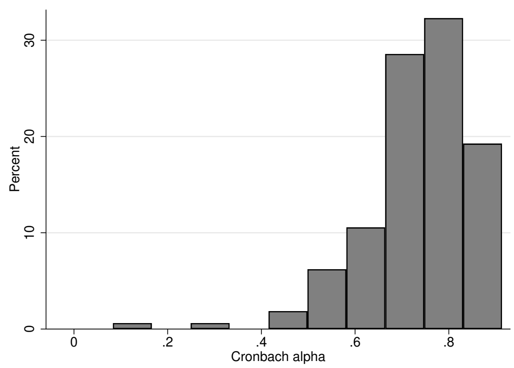

(b) Question 1

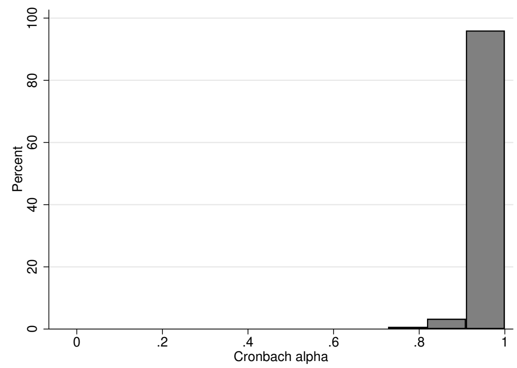

(c) Question 2

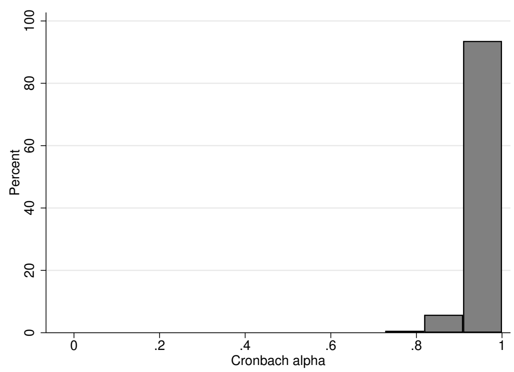

(d) Question 3

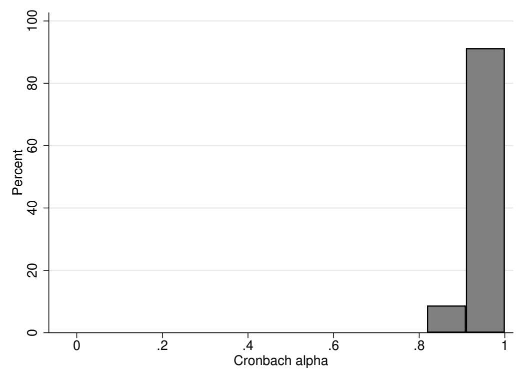

Table 23: Effect of salary on saves, controlling for perceptions

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

Table 24: Effect of salary on views, controlling for perceptions

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