

Matching Frictions and Distorted Beliefs: Evidence from a Job Fair Experiment *

Girum Abebe,[†] Stefano Caria,[‡] Marcel Fafchamps,[§] Paolo Falco,[¶]
Simon Franklin,^{||} Simon Quinn^{**} and Forhad Shilpi^{††}

October 13, 2022

Abstract

We evaluate the impacts of a randomized job-fair intervention in which jobseekers and employers can meet at low cost. The intervention generates few hires, but it lowers participants' expectations and causes both firms and workers to invest more in search as predicted by a theoretical model; this improves employment outcomes for less educated jobseekers. Through a unique two-sided belief-elicitation survey, we confirm that firms and jobseekers have over-optimistic expectations about the market. This suggests that, beyond slowing down matching, search frictions have a second understudied cost: they entrench inaccurate beliefs, further distorting search strategies and labour-market outcomes.

JEL codes: O18, J22, J24 , J61, J64.

Keywords: job-search strategy; recruitment; matching; expectations; beliefs; reservation wage; youth unemployment; Ethiopia.

*We thank audiences at the CSAE Annual Conference, the IZA/DFID GLM-LIC Conference, RES Annual Conference, SEEDEC Conference, SERC Work in Progress Seminar, and NEUDC; in particular, Prithwiraj Choudhury and Pieter Serneels who provided helpful comments and feedback. We thank Jali Bekele, Biruk Tekle, Marc Witte, Alemayehu Woldu and Ibrahim Worku for outstanding research assistance. Data collection and experimental implementation were funded by The World Bank Research Support Budget (RSB), GLM-LIC, and by the International Growth Centre. The views expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the IBRD/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent. The project would not have been possible without the constant support of Rose Page and the Centre for the Study of African Economies (University of Oxford), nor without the support of EconInsights in Addis Ababa. This RCT was registered in the AEA Registry for RCTs under Trial number AEARCTR-0000911.

[†]World Bank, Addis Ababa: gtefera@worldbank.org

[‡]University of Warwick: stefano.caria@warwick.ac.uk

[§]Freeman Spogli Institute, Stanford University: fafchamp@stanford.edu

[¶]University of Copenhagen: paolo.falco@econ.ku.dk

^{||}Queen Mary University of London: s.franklin@qmul.ac.uk

^{**}University of Oxford: simon.quinn@economics.ox.ac.uk

^{††}World Bank: fshilpi@worldbank.org

1 A matching experiment

Matching frictions can prevent the efficient allocation of workers across firms, sectors and industries. The costs of this misallocation are especially large in developing countries, where profound economic transformations are causing rapid changes in the labour market (Bryan and Morten, 2019; Hsieh, Hurst, Jones, and Klenow, 2019). However, the precise channels through which matching frictions make it difficult for firms to hire the right workers – or for workers to find the right firm – are not well understood.

In this paper, we provide new experimental evidence on the distortions generated by matching frictions in a fast-growing developing economy (Alfonsi, Bandiera, Bassi, Burgess, Rasul, Sulaiman, and Vitali, 2020; Abebe, Caria, Fafchamps, Falco, Franklin, and Quinn, 2021). In particular, we focus on the high costs that firms and workers have to bear to meet each other. These costs have well-known direct effects on the labour market: they slow down hiring and job-finding, and reduce the likelihood of good matches between workers and firms.¹ They can also have an understudied *indirect* effect. By reducing the exposure of firms to workers and of workers to firms, these costs can entrench inaccurate beliefs about the labour market, which, in turn, are likely to distort job-search and recruitment decisions, resulting in fewer and poorer matches. Our central contribution is to shed light on this second channel – through an experiment that provides a large one-time reduction in the cost of worker-firm meetings, and through a novel survey that provides corroborating evidence on workers and firms’ misperceptions about the labour market.

Our evidence comes from a country that is undergoing a rapid economic transformation: Ethiopia. Similar to many other fast-growing economies in Sub-Saharan Africa, Ethiopia is witnessing the expansion of non-traditional economic sectors, sustained workforce growth, and a strong build-up of secondary education. In this context, updating beliefs about the changing market fundamentals is particularly important. At the same time, acquiring information can be costly for workers and firms (Abebe, Caria, and Ortiz-Ospina, 2021).

We evaluate the impacts of reducing these costs by inviting to a job fair a sample of young jobseekers and of medium-to-large formal employers. At the fair, workers can meet several employers at a low marginal cost, and employers can easily talk to many potential

¹ Eeckhout (2018) summarises some of the key theoretical literature. Additionally, recent structural work in both developed and developing countries consistently detects the presence of meaningful labour market search costs (DellaVigna, Lindner, Reizer, and Schmieder, 2017; Van den Berg and van der Klaauw, 2019; Abebe, Caria, and Ortiz-Ospina, 2021).

young recruits. As a result, each side of the market has the opportunity to gather a large amount of information about the other side at very low cost. Our experimental design is uniquely placed to fully explore matching frictions on both sides of the market. First, we randomize participation within representative samples of both firms and unemployed jobseekers. Second, we collect data on their baseline search strategies, including their labour-market expectations. Third, we gather rich information on their interactions at the fair, including meetings, interviews, and job-offers. Finally, we observe workers and firms' search strategies over the months following the fair. This allows us to test for immediate as well as delayed treatment effects on firms and workers.

We show that the fairs generate few immediate hires (one for every 12 firms that attended), but they lead to a change in search strategies and an increase in search effort among both firms and jobseekers. Specifically, after the fairs, firms increase the amount they spend on advertising positions and jobseekers increase the frequency of their visits to job boards. Moreover, firms decrease their reservation quality for workers (shifting towards less-educated ones) and jobseekers reduce their reservation wages. The effects on jobseekers are concentrated among those without tertiary education, whose employment prospects improve significantly as a result (permanent employment rates double and formal employment rates increase by almost 50 percent).

This evidence suggests that the job fairs impacted firms' and job-seekers' search-strategies by correcting initial over-optimism. To validate this hypothesis, we first provide evidence of over-optimism in our sample at baseline. Second, we provide evidence that the job fairs moderate this over-optimism on both sides of the market: treated firms expect to hire fewer workers and treated workers set more realistic reservation wages. Third, we use a theoretical model to show how an information shock can correct beliefs and spur higher search effort, in line with our results.

In the final part of the paper, we delve deeper into the inaccurate beliefs that firms and jobseekers bring to the labor market. Specifically, after we analysed the effects of the job fair, we went back to the field and ran a unique survey administered to a new sample of firms and young jobseekers – each side selected respectively during hiring and during job search. Since we simultaneously observe representative samples of both sides of the market, this survey enables us to contrast beliefs with actual data on the true distributions. Our aim was twofold. First, we wanted to corroborate our findings on over-optimism. Experiments are typically not suitable to indisputably establish over-optimism, since they rarely have access to representative samples of both workers and firms (and hence what appears to be a mismatched belief in the experimental data may

be the result of selection into the experimental sample, and not a genuine misperception). Second, while the previous literature has mostly focused on individuals' forecasts of their own future outcomes, we wanted to explore the beliefs that individuals hold about the structural features of the labor market { e.g. wage or ability distributions. These beliefs are central to the matching process, but have received limited attention in the literature.

The additional belief-elicitation exercise shows clear evidence that firms and jobseekers hold incorrect beliefs about labor-market fundamentals. Firms overestimate the ability of jobseekers, while jobseekers underestimate how difficult it is to attain higher-paid professional jobs. On the other hand, we show that jobseekers have relatively accurate beliefs about facts that are easy to observe (e.g., the distribution of wages across occupations). This suggests that while market participants value and acquire relevant publicly available information, the natural interaction between firms and young jobseekers in the labour market is insufficient for participants to learn key parameters of the matching process that are observed through direct interactions (e.g., the ability of workers). Crucially, this is consistent with the hypothesis that our experiment helped to correct underlying misperceptions.

Our paper contributes to a small but growing literature on matching frictions in developing countries' labor markets. First, we are the first to provide evidence that the search strategies of both firms and jobseekers are based on incorrect beliefs. Our results show that both firms and jobseekers revise these strategies in response to an intervention that allows each side of the market to acquire more information on the other side. This is a novel finding in a literature that has so far been especially silent on the search decisions of firms. Two recent papers have begun to fill this gap by highlighting the role of recruitment costs and uncertainty about candidate ability (Singh et al., 2022; Hensel et al., 2022). Consistent with their findings, several experiments have documented that skill certification can improve jobseekers' employment prospects (Abebe et al., 2021; Bassi and Nansamba, 2021; Carranza et al., 2021). But no study to date has shown that firms' recruitment strategies respond to interventions that increase exposure to jobseekers.² This finding is surprising, as firms are often assumed to have rational, accurate expectations. It also suggests that the distortions generated by workers' miscalibrated search strategies are unlikely to be offset by optimal search strategies set by firms. Matching frictions may thus be deeper than previously thought.

² Our findings complement the results in Abebe et al. (2021), who show that firm managers make biased forecasts about the results a recruitment intervention, and of Caria and Falco (2021), who present lab-experimental evidence that small-firm managers are excessively concerned about worker trustworthiness.

Second, we show that young jobseekers are overly optimistic about their employment prospects, and that moderating optimism can boost search effort and improve employment outcomes³. In contrast, recent studies in South Africa and Uganda document interventions that reduce over-optimism, but also depress job search and job finding (Banerjee and Sequeira, 2021; Bandiera et al., 2021). The model we propose helps reconcile these findings, since it shows the conditions under which correcting overoptimism will boost or depress job search effort. Correcting over-optimism is a particularly attractive policy option in settings where it leads to more search effort and employment, as there are positive social benefits from expanding young individuals' labor market participation in low income countries. Additionally, interventions that lead to more search can enable individuals to acquire useful information on other parts of the job search process, and hence correct a wider set of beliefs than those originally affected by the policy.

Finally, we expand our understanding of the inaccurate beliefs that jobseekers and firms bring to the labor market. This is important, as the drivers of over-optimism in labor markets are not well understood. Are jobseekers and firms overconfident about their own prospects relative to those of their peers, or do they instead misunderstand the fundamental parameters of the labor market in which they operate? Our detailed follow up survey shows that representative samples of both firms and workers hold inaccurate beliefs about the fundamentals of the labour market. As discussed above, such misperceptions have important consequences, as these beliefs are central to the choice of an effective search strategy. By contrast, models of search and matching typically assume that, while market participants may suffer from a problem of asymmetric information | that is, they may be uncertain about the skills of a particular worker or the characteristics of a particular job | they have accurate information about relevant distributions from which these variables are drawn (Rogerson et al., 2005; Terviö, 2009; Wright et al., 2021). In this paper, we show that firms and jobseekers make decisions on the basis of information that is limited in a more fundamental sense.

2 The study population

We work in a rapidly growing urban center in a low-income setting, where frictions are likely to be prevalent in the labor market. Addis Ababa, capital of Ethiopia, is a good choice because it combines these characteristics with the additional feature that, at the

³ These findings are consistent with a recent literature studying the relationship between optimism and job finding in rich economies (Spinnewijn, 2015; Krueger and Mueller, 2016; Mueller et al., 2021; van der Klaauw and Ziegler, 2021).

time of our study, the main avenue through which firms advertise openings is through job-vacancy boards located in the center of the city. While the purpose of these boards is to facilitate job search, they nonetheless entail sizeable transaction costs { especially for jobseekers, who must incur substantial transport costs to visit, and then need to spend considerable time visually scanning the boards to identify suitable openings.

Screening by firms is also challenging, given the limited information that can be extracted from the CVs of young labor market entrants (Abebe et al., 2021). Like many growing cities in the developing world, Addis Ababa has recently experienced a large increase in the number of available jobs, coupled with high in-migration flows. This makes it hard for firms and jobseekers to have accurate beliefs about the distribution of wages, employment opportunities, and workers' abilities. All these features suggest that job fairs are a promising intervention in this context.

2.1 Surveying jobseekers

The job fair intervention reported in this paper draws on the same sampling frame as Abebe et al. (2021) and was partially run alongside that experiment⁴. The study involves a representative sample of young educated jobseekers in Addis Ababa. To select our sample, we first define geographic clusters using enumeration areas from the Ethiopian Central Statistical Agency (CSA).⁵ Our sampling frame excludes clusters within 2.5 kilometres of the center of Addis Ababa and clusters outside the city boundaries. Clusters are selected at random from the sampling frame. To minimize potential spillover effects across clusters, we impose the condition that directly adjacent clusters cannot be selected together.

In each selected cluster, we used door-to-door sampling to construct a list of all individuals who: (i) are aged between 18 and 29 (inclusive); (ii) have completed high school; (iii) are available to start working in the next three months; and (iv) are not currently working in a permanent job or enrolled in full time education. We randomly sample individuals from this list to be included in the study. The lists include individuals with different levels of education. We over-sample individuals with post-secondary education to ensure that they are sufficiently represented in our sample.

All randomly selected individuals were contacted to establish their willingness to participate in the study and be interviewed. We completed baseline interviews with 4,388

⁴ Abebe et al. (2021) report two parallel field experiments: a transport subsidy to visit job boards, and a workshop intervention to help jobseekers to signal their cognitive and non-cognitive skills to employers.

⁵ CSA defines enumeration areas as small, non-overlapping geographical areas. In urban centers, these typically consist of 150 to 200 housing units.

eligible respondents. We attempted to contact individuals by phone for at least a month (three months on average) and dropped individuals who could not be reached after at least three attempts. We also dropped any individual who had found a permanent job at the time of baseline and had been in this job at least six weeks. Finally, we dropped individuals who had migrated away from Addis Ababa during the phone survey. In all we were left with 4,059 individuals included in our experimental study. Of these 1006 were invited to the jobs fairs. Another 2226 were involved in the experimental interventions discussed in [Abebe et al. \(2021\)](#), while 823 remain in the control group.

We collected data through both face-to-face and phone interviews. We completed baseline face-to-face interviews between May and July 2014 and endline interviews between June and August 2015. Information was collected on the socio-demographic characteristics of study participants, their education, work history, nances, and their expectations and attitudes. We also kept in touch with all study participants by phone throughout the duration of the study, at which time we administered a short questionnaire on job search and employment⁶.

We have low attrition: 93.3% of baseline respondents were re-interviewed at endline. Few covariates predict attrition and we are unable to reject a joint F -test that a set of key covariates have no effect on attrition (see Appendix Table [B.10](#) in the Online Appendix). However, we do find that the individuals invited to the job fairs are slightly more likely to respond to the endline survey. Yet, because attrition is low overall (8% in the control group and 5.6% in the treatment group), we are not concerned that this affects our main results. Our key findings are robust to bounding our estimates using the method of [Lee \(2009\)](#). Attrition in the phone survey is also low; for example, we were still able to contact 90% of the respondents in the final month of the study⁷.

2.2 Surveying firms

We surveyed 498 large firms in Addis Ababa. These firms were sampled so as to be representative of large employers in the city, stratified by sector. All major sectors in the economy are covered, including construction, manufacturing, banking and financial services, hotels and hospitality, and other professional services. To sample firms, we first compiled a list of the largest 2,178 firms in Addis Ababa. Since no firm census exists for

⁶ [Franklin \(2018\)](#) shows that high-frequency phone surveys of this type do not generate Hawthorne effects and do not affect jobseekers' responses at endline.

⁷ Appendix Figure [B.1](#) shows the trajectory of monthly attrition rates over the course of the phone survey.

Ethiopia, we rely on a variety of data sources, including lists of formal firms maintained by different government ministries. In all, we gathered data from more than eight different sources. For the manufacturing sector, we rely on a representative sample of large firms that took part in the Large and Medium Enterprise surveys conducted by the Central Statistics Agency (CSA). For other sectors we requested lists of the largest firms from the government agency in charge of that sector. Whenever information on firm size is available, we impose a minimum size cut-off of 40 workers.

We draw the firms in our sample using sector-level weights to reflect the number of employers in that sector in the city. We construct these weights using representative labour-force data.⁸ The firms are, on average, large by Ethiopian and African standards. The mean number of employees per firm is 171.5. This masks considerable heterogeneity, particularly in the 'Tours & Hospitality' sector which is dominated by small hotels and restaurants; when this sector is excluded, average firm size is 326 workers. Detailed information on firms' total employment is given in Table 1, excluding casual daily laborers. On average, firms report employing 34 casual laborers per day.

The firms in our sample are growing in size and looking to hire new workers. At baseline, the median number of workers that a firm expects to hire in the next 12 months amounts to 12% of its current workforce. The median rate of hiring is highest (16%) among service sector firms, which are also the most likely to come to the job fairs. The most common types of workers whom firms expect to hire are white-collar workers, usually requiring university degrees. For details, see Appendix Table B.5.

2.3 Evidence of mismatched expectations at baseline

Our baseline data shows clear evidence of mismatched expectations between workers and firms { particularly in the market for lower educated jobseekers. On one side of the market, the reservation wages and future earning expectations of low-educated workers are significantly higher than what firms are prepared to pay. On the other side, firms appear overoptimistic about the prospect of filling positions with low-education workers.

Table 2 presents summary statistics on this mismatch. Panel A compares jobseekers' reservation wages with current wages at the firms that were invited to the fair. We show that jobseekers with a high-school diploma and no experience report a median reservation

⁸ Table B.7 in the Online Appendix shows the number of firms surveyed in our sample, divided into five main categories. Column (2) provides weighted percentages obtained by applying the inverse of the weights used to sample the firms. For instance we surveyed NGOs ('Education, Health, Aid') relatively infrequently because of the large number of NGOs in the data.

wage of 1,300 Birr per month. This is significantly higher than the median salary of 855 Birr that firms report offering to high-school graduates with no work experience. In contrast, jobseekers with tertiary education have reservation wages that are well within the boundaries of what is available in the market: on average, firms report paying recruits with a university degree around 4,500 Birr per month, which is well above the median reservation wage of 2,500 Birr reported by university graduates without experience at the fairs. Indeed, only 10% of tertiary graduates in our sample have a reservation wage above the average wage paid for employees with their qualifications.

We also find a mismatch when we compare jobseekers' expectations of future earnings at baseline with actual earnings at endline (Table 2, Panels B and C). At baseline, the median response to the question "what will the average wage for someone of your skills and education be in one year?" is 1500. By contrast, the median wage among low-educated participants at endline is 1000. Similar mismatches exist over longer horizons: at baseline, the median expected wage in five years is 3,500. In our second endline (roughly five years after the intervention), median earnings are only 1800. The mismatch is largely driven by individuals without previous work experience (the vast majority of the sample), suggesting that lack of exposure to the labour market is driving these unrealistic wage expectations.

Finally, jobseekers are also overly optimistic about the number of job offers they will receive and the probability of having a job in the future. At baseline, the average expected number job offers over the next 4 months is 1.3. At endline, the actual average total number of offers received over the past 12 months is only 0.42. At baseline, 70% of high-school graduates say there is a "medium", "high" or "very high" probability, (30% say "high" or "very high") that they will find a job in the next 12 months. However, at endline only 50% have any job and only 6% have a permanent job (Table 2, Panel C).

On the firm side, wages offered are clearly below what workers are willing to work for, at least for low educated workers, suggesting that they may find it more difficult to hire than they expect. We find some evidence of hiring challenges among these firms. At endline, we find that firms in our control group have on average 2.1 unfilled vacancies open, while they hired only over 3.9 workers over the last year.

⁹ See the control means in Table B.12.

3 Experimental design

3.1 Randomization to the job fair

We assign treatment of jobseekers to the job fairs by geographical cluster, after blocking on cluster characteristics (see [Abebe et al. \(2021\)](#) for further details). The sample is balanced across all treatment and control groups, and across a wide range of outcomes { including baseline outcomes that are not used in the stratified randomization procedure. We present extensive balance tests in Appendix Table [B.1](#). For each baseline outcome of interest, we report the p -values for a test of the null hypothesis that we have balance between treatment and control groups. We cannot reject the null for any of variables that we study.

We assign firms to either a treatment group or a control group. Firms in the treatment group are invited to attend the job fairs; control firms are not. We use block-level randomization techniques suggested by [Bruhn and McKenzie \(2009\)](#). Firms are first partitioned into five main industries (see Appendix Table [B.7](#)). Within each industry, firms are partitioned into blocks of four nearest neighbors on the basis of their Mahalanobis distance over a set of baseline variables¹⁰. We then randomize the firms in each block into two groups of two firms: one firm is invited to the first day of the job fair; the second is invited to the second day (see below for details); and the other two are assigned to serve as controls. Given the relatively small size of the firm sample, we use a re-randomization approach to ensure balance on a set of baseline covariates listed in Table [B.2](#).

3.2 Implementation of the job fairs

We invited treated jobseekers and treated firms to attend two job fairs. The first fair took place on October 25 and 26, 2014. The second fair took place on February 14 and 15, 2015. We run two fairs to increase the chance that each jobseeker and firm is able to participate in at least one of them. The job fairs were held at the Addis Ababa University campus, a central and well-known location. To minimize congestion, each job fair lasted two days and a randomly selected half of the firms and jobseekers were invited to attend on each

¹⁰ The variables used for blocking are listed in Appendix Table [B.8](#).

¹¹ Following the recommendations of [Bruhn and McKenzie \(2009\)](#), we control for these covariates in our estimation, as well as for the baseline covariates used to construct the randomization blocks. Details of these variables and how they are defined are contained in our detailed pre-analysis plan. Simulations show that, with this sampling strategy, we have 78% power to detect a small treatment effect of 0.2 standard deviations at a significance level of 0.05%.

day. The firms that were invited to attend on Saturday 25 October were then invited to attend on Sunday 15 February; firms invited for Sunday 26 October were invited for Saturday 14 February. In contrast, jobseekers invited to attend on the Saturday of the first fair were also invited to attend on the Saturday of the second fair; jobseekers invited for the Sunday of the first fair were invited for the Sunday of the second fair. This ensures that, in each job fair, jobseekers are exposed to a different pool of firms, and that firms are exposed to a different pool of jobseekers.¹²

During each fair, jobseekers and firms are free to interact as they see fit. Each firm sets up a stall before the jobseekers arrive. These stalls are typically staffed by the firm's HR team who bring with them printed material advertising the firm. In a typical interaction, a jobseeker approaches the stall of a firm and asks questions about the firm and its vacancies. The firm's HR staff is then free to check his or her CV and to ask about the jobseeker's skills and work experience. If the jobseeker looks suitable for one of the firm's vacancies, the firm invites her or him for a formal job interview a few days after the job fair.

To avoid self-selection out of the sampling frame, we do not restrict invitations to the fairs to currently unemployed jobseekers, or to firms that have open vacancies at the time of the fair. Of our initial sample of jobseekers, only about 8% had permanent jobs by the time of the first job fair, and thus most jobseekers were still searching for work. Similarly, most firms were hiring at the time that the job fairs were held. 89% hired at least one worker in the year of the study and, on average, firms hired four workers in the month after the job fairs and 52 workers in the year of the fairs.

In total, we invited 1,007 jobseekers and 248 firms to attend the fairs. Both jobseekers and firms were contacted by phone, were given some information about the nature of the fairs, and had the opportunity to ask questions. Among firms, 170 attend at least one job fair, which represents quite a successful take-up rate of 68.5%. Of the firms that do not attend the fairs, 12% say it is because they do not have an open vacancy at the time. The remaining firms tend to cite logistical issues or previous commitments. Only 13 firms respond that they would not find the job fair useful.¹³

¹² Weekend days are selected to maximize the opportunity for both firms and jobseekers to attend. In preliminary discussions with firms, we realized that most would be unable to take time off from daily activities to attend during the week, but they were interested in attending on a weekend. Similarly, many jobseekers in our sample work in casual jobs and are more likely to be unavailable during the week. Since many Ethiopians attend religious services on the weekend, we set a long enough time window for jobseekers to be able to attend.

¹³ In Appendix Table B.9, we run a descriptive regression to explore correlates of firm attendance at job fairs.

Of the 1007 invited jobseekers, 606 attend at least one fair, a 60% take-up rate. The most common reason that jobseekers give for not attending the fairs is that they are busy during that particular weekend. This reason is given by 226 jobseekers in the first fair and 229 jobseekers in the second. Other reasons include not being able to take a new job (9 jobseekers at the first fair and 83 at the second) and finding the venue of the fair hard to reach (31 respondents for the first fair and 25 for the second).

Two baseline characteristics predict higher attendance by jobseekers: search effort at baseline; and whether the jobseeker uses a school certificate during job search. It follows that jobseekers who attend the fairs are, if anything, more active and organized in their job search. Those who attend are also more likely to have a university degree or diploma, but this is not statistically significant. Taken together, this evidence provides reassurance that results are not driven by negative selection of jobseekers coming to the fairs.

3.3 Matching at the fairs

At the beginning of each fair, we give jobseekers a list of all the firms invited, with basic information on the firm. In the second fair, we also give jobseekers the list of all vacancies. We give firms a list of all jobseekers invited to the fairs, with some information about their education and past work experience. We then ask firms to list up to 10 jobseekers with whom they would like to talk at the job fair. After collecting the list of requested meetings from each firm, we post them on a board at the fair¹⁴.

In order to increase match efficiency and avoid congestion at the fairs, we create a list of 15 recommended meetings that we give to each jobseeker at the beginning of the fair. Of the 15 firms on the recommended list, 10 are selected using a Gale-Shapley Deferred Acceptance algorithm described below ([Gale and Shapley, 1962](#)); the other five are selected at random. The order of presentation on each list is similarly randomized. We tell jobseekers that these are the firms they should talk to during the fair. Each firm similarly receives a personalized list showing the names of all the jobseekers who have been recommended to meet that firm. The recommendations are based on information about firms' vacancies that we obtain through a phone survey shortly before the fair.

The purpose of the Gale-Shapley algorithm is to suggest sensible matches for these vacancies, given baseline characteristics of both jobseekers and firms. Indeed, not all jobseekers are qualified for certain positions, and not all firms can attract the best jobseekers.

¹⁴ Given the logistics of collecting lists of names from more than a hundred employers, the lists were posted a few hours after the beginning of the fair.

To avoid firms and jobseekers wasting time in unsuccessful meetings, we seek to pair those firms and jobseekers who, given the distribution of firms and jobseekers at the fair, stand a higher chance of leading to a hire. To this end, we start by constructing a synthetic ranking of all vacant positions for each jobseeker, and similarly a synthetic ranking of all jobseekers for each firm. The rankings of jobseekers by firms are constructed using lexicographic preferences over: (i) whether the jobseeker held a previous occupation that matches that of the vacancy; (ii) the jobseeker's educational qualification for the job; and (iii) the jobseeker's years of wage employment. The rankings of jobseekers vary across firms. For the jobseeker's rankings of vacancies, we use a simple ranking over the advertised wage. This means that, for the purposes of forming recommendations, all jobseekers synthetically rank vacancies in the same way.

These rankings are not intended to represent literally the true preferences of all participants over all possible matches. Indeed, gathering information on all these preferences would have been logistically impossible in the allotted time and any attempt to impose such a ranking burden on jobseekers or firms would undoubtedly have reduced substantially the participation in the experiment. Rather, the rankings are intended to provide a fast way of improving on random encounters at the fairs that takes into account the heterogeneous set of vacancies and jobseeker skills that are present at the fair. After creating a ranking of jobseekers for each vacancy and a ranking of vacancies for each jobseeker, a Gale-Shapley algorithm is used to match jobseekers and firms. Specifically, the algorithm generates a single set of matches; we then iterate the algorithm 10 times, requiring a different set of recommended matches each time¹⁵. This generates the 10 recommended matches mentioned above; to this list, we then add five random matches.

Figure 1 illustrates the outcome of the matching algorithm. Each point represents a stable match recommended by the algorithm. The figure shows which combinations of firm rankings and jobseeker rankings generated these recommended matches. The graph provides a visual illustration that the algorithm worked well (in the sense of generating matches between firms and jobseekers who are, on the basis of jobseeker skills and experience, reasonably suitable for each other). Note the substantial mass at the bottom-left of the graph; this mass shows that, for those firms paying higher wages, the algorithm recommends matches that provide a reasonable occupational fit. For example, for the top 100 firms in the jobseekers' ranking, the median match is to a jobseeker with a firm ranking of just 14, that is, a jobseeker ranked quite high according to our synthetic firm preferences.

¹⁵ We implement this requirement by taking the matches recommended in iteration t and placing those matches at the bottom of the firms' and jobseekers' rankings in iterations $s > t$.

4 The effects of the fairs

In this section, we document the impacts of the job fairs on employment outcomes and search behaviour { both at the level of the jobseeker and at the level of the firm. We begin by presenting evidence on hiring that took place at the fairs and in their immediate aftermath. We then present impacts on employment and search outcomes at endline (six months after the second fair), as well as impacts on jobseekers' and firms' expectations. We also present evidence on the trajectories of impacts based on a high-frequency survey conducted between baseline and endline.

We measure employment outcomes through data on job interviews, offers, hiring, and employment in different types of jobs. To test for impacts on search behaviour, we estimate treatment effects on the search strategies used by firms and jobseekers at endline (e.g., amount spent on advertising vacancies, amount of time spent looking for jobs). It should be noted that impacts on search behaviour are more likely to be observed if a direct effect on hiring is absent or weak: if treated jobseekers find a job and firms fill their vacancies as a direct outcome of the fairs, they have little cause to revise their expectations and search strategy.¹⁶ Finally, we test for impacts on jobseekers' and firms' expectations at endline by analysing the effect of the treatment on jobseekers' reservation wages and on firms' hiring expectations.

4.1 Immediate impacts: Hiring at the job fairs

The fairs generated rich interactions between firms and jobseekers. 454 jobseekers (75% of those attending) interacted with at least one firm at the job fair, either through an informal interview or an in-depth discussion with a recruiter. This finding is particularly strong among participants who benefited from the matching algorithm treatment (as discussed below). In total, we record 2,191 contacts between firms and jobseekers.

The interactions at the fairs resulted in 105 formal job interviews conducted at participating firms in the immediate aftermath (the finding is based on a phone survey conducted immediately after each fair).¹⁷ Further, these 105 interviews are concentrated on 67 jobseekers only, representing 11% of those attending the fairs. These interviews led

¹⁶ The pre-analysis plan that we led for this experiment can be found at <https://www.socialscisceregistry.org/trials/1495>. Most pre-specified outcome families are presented in the order in which they appear in the pre-analysis plan. Those that are not documented here in detail can be found in the Online Appendix.

¹⁷ This implies a rather low conversion rate of 1 interview for each 20.9 contacts established at the fair (in the open market, we estimate that jobseekers get an interview every 3.5 applications made).

to 76 offers (made in the immediate aftermath of the fairs) to a total of 45 jobseekers, which represents a healthy conversion rate of one offer for each 1.4 interviews (and compares favourably with the open-market rate of 1.9). Contrary to what one might expect in a job fair for educated jobseekers, offers were disproportionately made to less-educated applicants.¹⁸ A large majority (81%) of offers, however, were rejected. The offer rejection rate is particularly high among less educated jobseekers: 85% for applicants with a high-school diploma compared with 71% for those with tertiary education. Only 33% of offer recipients with a high-school diploma accepted one of their offers. We view these findings as prima facie evidence of a mismatch between workers' expectations and what firms are willing to offer. Overall, we find that the fairs had little immediate impact on hiring by treated firms (Appendix Table B.12 and B.13).

In Appendix A, we test whether the limited impacts of the fairs on hiring may be due to the market being too thin (too few high-quality matches available), or to problems of congestion and mis-coordination during the fairs. We have evidence against both hypotheses. First, we document that the firms attending the fairs had a large number of open vacancies at the time of the event and the occupational composition of those vacancies exhibits considerable overlap with the distribution of occupations desired by invited jobseekers. Second, using dyadic data on firm-worker interactions, we show that our stylized matching algorithm was useful in identifying matches that were deemed worth pursuing by market participants. Specifically, we show that our synthetic rankings strongly predict both requested meetings and actual meetings. The fairs thus appear to have reached their objective of facilitating meetings between jobseekers and the firms that suited them best. This reassures us that our setup managed to minimise wasteful interactions and potential congestion.

4.2 Endline impacts on optimism, search, and employment

In this section, we examine the impacts of the intervention at endline (six months after the second fair). We report impacts on firms and workers separately. For the latter, we also leverage a high-frequency survey to study the trajectories of impacts in the months following the job fairs. Overall, we find clear evidence that both firms and jobseekers increase their search effort as a result of being invited to the fairs, and this leads to changes in employment outcomes that are particularly evident for the group of jobseekers that revised their search strategies the most (less-educated workers).

¹⁸ 55 offers (72%) went to jobseekers with at most a high-school diploma, even though they represented a minority of the jobseekers attending.

4.2.1 Impacts on firms' search strategies and hiring outcomes

First, we study the impact of the intervention on the search and hiring outcomes of firms measured at endline. To this end, we use an ITT approach with an ANCOVA specification. Following current practice, covariates used for balancing the randomization are included as controls. For each outcome of interest, we estimate regressions of the form:

$$y_i = \alpha_0 + \alpha_1 \text{fairs}_i + y_{i;\text{pre}} + x_{i0} + \epsilon_i; \quad (1)$$

with robust standard errors. Variable $y_{i;\text{pre}}$ is the dependent variable measured at baseline and x_{i0} includes the randomization variables listed in Table B.2. In the tables, we show each regression as a row and we report the estimated ITT $\hat{\alpha}_1$, the mean of the control group, and the number of observations. We report both p -values and False Discovery Rate q -values, the latter being calculated across the family of outcomes (Benjamini et al., 2006).¹⁹

Our first finding is that, as a result of the job fairs, firms invested more in worker search and recruitment. Our regression estimates, presented in Table 3, show that treated firms are six percentage points more likely to advertise new vacancies in the last 12 months, relative to a control mean of about 79%. They are also 12 percentage points more likely to advertise for professional positions, relative to a control mean of 60%. Firms are also almost 10 percentage points more likely to advertise their vacancies on the job boards, relative to a control mean of 33 percent. All three effects are statistically significant after controlling for multiple hypothesis testing. This suggests that the intervention leads both firms and jobseekers to search more intensely through the main channels available to them at the time of the study.²¹

Our second finding is that firms reorient their hiring away from highly educated

¹⁹ Throughout this paper, we report the average treatment effect of the job fairs. As outlined in the pre-analysis plan, our study was designed to enable us to estimate separately an effect of the fairs both with and without additional information revelation about workers' abilities. Since we found no direct effect of the fairs on hiring for either treatment arm, and the experimental information revelation was designed specifically to improve direct hiring at the fairs, we took the decision to pool the treatment arms. This improves the precision of our null estimates of the direct effects of the fairs.

²⁰ Throughout the analysis, we distinguish between professional workers and non-professional workers. 'Professional workers' refers to traditional notions of 'white-collar employees': typically those with some degree or diploma working in relatively highly skilled positions. For manufacturing firms, 'non-professional workers' refers mostly to production workers; for service-based firms, these include mostly workers dedicated to client services (tellers, waiters, receptionists etc).

²¹ We do not find significant heterogeneity in these impacts. However, effects on recruitment appear to be larger among firms that did not hire many young people at baseline (Table B.20). This result is consistent with a learning story, as those firms likely have noisier and more inaccurate priors about young people in the labour market.

workers, in particular for professional positions. In other words, they lower the reservation quality of their hires. Among firms that hire above the median number of professional workers (a pre-specified dimension of heterogeneity) we find that, beyond raising recruitment investments, the firms also (i) significantly reduce the proportion of professional workers with degrees by about 7 percentage points (relative to a control mean of 72 percent), (ii) reduce hiring by an average of 17 workers (over a control mean of 62 workers), and (iii) reduce overall firm size.²² These results are consistent with the fact that firms at the job fairs did not extend interviews to workers with degrees (as discussed above). Importantly, they are also consistent with the finding (discussed in the next section) that firms in this labour market have overly optimistic beliefs about the ability premium of highly educated workers. The firms lowered their expectations. Another piece of evidence consistent with this conclusion is that at endline treated firms report lower expectations of hiring workers for occupations that require higher educational levels over the next twelve months (Table 6).

The changes in firms' recruitment strategies have only small effects on aggregate hiring outcomes at endline. We find a small but significant increase in unfilled vacancies over the 12 month period from baseline to endline (in Panel A of Table B.11). We find no impact on the time taken to fill open positions or on firms' reported costs of recruitment. We find no significant impact on the number of people hired in the last 12 months, the hiring of job candidates with a degree, or hiring on a permanent contract (Panel B of Table B.11).²³ This is perhaps not surprising, given that our sample is composed of large firms that hire on average 56 new workers per year. In these firms, changes in recruitment practices may take a longer time to affect the overall composition of the workforce.

4.2.2 Jobseekers' search strategies and employment outcomes

Next, we examine the effect of the treatment on jobseekers' search and employment outcomes. The specification we estimate is the same as equation 1, but we now focus on jobseekers rather than firms. To account for the fact that jobseekers were randomized to treatment according to their enumeration area of residence, standard errors are clustered by enumeration area. We report both conventional p -values and False Discovery Rate q -values.

We find that the treatment increases jobseekers' search effort. This is reflected

²² Results (i) and (ii) are presented in Table B.24 in the Appendix, result (iii) is shown in Table B.23.

²³ Similarly, we find no impact on the firms' overall workforce composition (Appendix Table B.14), overall turnover and employee growth (Appendix Table B.15), and general HR practices (Appendix Table B.16).

in a higher number of visits to job boards at endline as shown in Panel A of Table 4: treated jobseekers report roughly three more visits to the job boards, relative to a control mean of 15. We can plot the trajectory of the effects on search over time, using high-frequency phone call survey conducted between baseline and endline. Figure 2 shows that the probability that a jobseeker visits the boards goes up by about 8.3 percentage points (26 percent) in six weeks following the first job fair²⁴. Since the job boards are the main source of information on vacancies, this represents a sizeable increase in their effort to search for employment.

Turning to the impacts of the fairs on the employment outcomes of jobseekers, we find that the effects are concentrated among the least educated ones, who experience a large increase in employment quality due to the intervention. In Panel B of Table 4 we disaggregate treatment effects by whether or not the respondent has more than secondary education. Among the less educated jobseekers, we document an increase of 6 percentage points in the probability of having a permanent job relative to a control mean of just 6 percent at endline { i.e., a doubling of the probability of permanent employment. We similarly find an increase in the probability of having a formal job by 5 percentage points relative to a control mean of about 11 percent { i.e., a 45% improvement²⁵. It is important to note that less educated workers are also the ones that experience the most significant changes in the search strategies, driving the average impacts discussed above (i.e., higher search effort and lower reservation wages²⁶). This is consistent with the hypothesis that changing search strategies in light of the information acquired at the fairs leads to better labour-market outcomes for jobseekers²⁷.

4.2.3 Experimental evidence on expectations

We find several key impacts that are consistent with the idea that the fairs affected search behaviour and eventually employment by providing a 'sobering' experience. For

²⁴ Specifically, we estimate this difference in probabilities using a Linear Probability Model in an AN-COVA specification, in which we regress job search on treatment, baseline search status and a sector of baseline balancing variables. We cluster at the level of individual jobseekers, and show both point estimates and 90% confidence intervals; we do this both by regressing on fortnight dummies, and by imposing a quadratic shape.

²⁵ In the bottom row of Table 4 we report p-values for the null hypothesis that treatment effects are equal across educational categories. The null is rejected for wage mismatch and having a permanent job, and it is close to being rejected ($p < 0.12$) for visiting job board and having a formal job.

²⁶ They reduce their reservation wages by 9 percent as a result of treatment, closing the mismatch between reservation wages and market wages by 7 percent. They increase their visit to the job boards by 4.2 percent, relative to a control mean of 11 visits.

²⁷ Additional treatment effects on employment outcomes, job amenities, and job search at endline are presented in Appendix Tables 4, B.3, and B.4.

jobseekers, Table 5 (Panel A, column 1) shows that treatment results in a significant 7 percent reduction in endline reservation wages. To test whether treatment brings reservation wages more in line with market conditions, we construct a 'wage mismatch' variable equal to the absolute difference between the log of the reported reservation wage, and the log of the average wage earned by a worker with the jobseeker's skill and education. We present treatment effects on this variable in column 2 (Panel A). We find that treatment reduces the wage mismatch by a significant 4 percent. Column 5 in Panel A also indicates a significant reduction in the probability that workers aspire to find a permanent job { which suggesting that they adjusted their aspirations downwards due to reduced expectations.

In Panel B of Table 5, we then split these impacts between respondents with only a high-school diploma and those with post-secondary education. We find that each of the impacts is driven by the less educated respondents { precisely the group showed, in section 2.3, to have a stark mismatch of expectations. Again, this is consistent with the hypothesis that the fairs had a sobering effect on jobseekers, which in turn caused them to increase their search effort.

For firms, we measured expectations of the number of workers that they would hire in the next year. The results are shown in Table 6, and are also consistent with the hypothesis that the job fairs had a sobering effect on expectations: specifically, we find that firms that attended the fair expect to hire approximately 30% fewer workers over the next 12 months, at the time of the endline survey.²⁸

5 A dynamic model of search under distorted beliefs

The results in section 4 show that the experience of attending the fairs { without having successfully found a match { persuaded each side of the market to revise their search strategies. On the one hand, firms hired very few jobseekers through the fairs, yet the fairs caused firms to invest more in worker search and recruitment. On the other hand, the fairs caused jobseekers to expend more effort on search, and to lower their reservation wages, and to be more likely to find jobs six months later. This immediately implies that each side of the market received an important shock to their beliefs.

In this section, we present a stylised theoretical framework that helps us to interpret

²⁸ Of course, for firms, it is intrinsically harder to find direct evidence of changes in overall hiring expectations. For many vacancies, new hires need to be found, and firms may have subtly downgraded their expectations of the quality of the candidates that they would find for each position.

the results outlined above, and offers an explanation for the observed changes in search behaviour after the fairs. Specifically, the model formalises the notion that firms and jobseekers held beliefs that made them overly optimistic about the possibility of finding a good match, given their existing investments in search. The model illustrates how a sobering information shock can cause an increase in search effort. We consider a firm trying to fill a single vacancy (we later discuss how the same framework also usefully captures the search problem of a jobseeker). The model incorporates important features observed among surveyed firms { notably, that firms (i) often have a specific notion of a minimum appointable standard when advertising a position, and (ii) seldom hire quickly for an advertised position.

The model firm searches in discrete time (with discount factor $\beta < 1$). In each period that the vacancy remains unfilled, the firm suffers a direct reduction in profit of $\delta > 0$; this could reflect, for example, the cost of being unable to proceed with a project for want of filling the vacancy. In each period that the vacancy remains unfilled, the firm may spend $s \geq 0$ to generate λ matches with prospective employees, such that $\lambda \sim \text{Poisson}(s)$. The quality of any given match $x > 0$ is drawn from some distribution $F_x(x; \mu)$, where the mean μ represents the firm's belief about the quality of available applicants. We denote by y the quality of the best match realised in a given period (where $y = 0$ if there are no matches); following [Bobotas and Koutras \(2019\)](#) and [Wilken \(2021\)](#), this best match has CDF $F_Y(y; s; \mu) = \exp[-F_x(y; \mu)] \lambda s$.

Having observed y , the firm decides whether to hire. The firm optimally does this using a cutoff rule, comparing y to some reservation match quality \bar{x} ; thus, the firm hires if $y \geq \bar{x}$ and otherwise prefers to leave the position open. We impose that the firm has some minimum match quality z , implied by the technical requirements of the position; this operates as a lower bound on the firm's choice of \bar{x} . For example, there is a minimum set of technical skills that a crane driver must reach before she can be employed { irrespective of how costly the firm finds it to leave the position vacant.

For simplicity (and following [McCall \(1970\)](#)) we assume that, if the firm hires, the contract is permanent { so that the value of meeting a best applicant with quality $y \geq \bar{x}$ is simply $V(y) = y/(1 - \beta)$. Note that the firm is indifferent between all values of $y \geq \bar{x}$; therefore, the value to the firm of leaving the position unfilled is defined recursively as:

$$V(0) = \max_{s \in [0, \bar{x} - z]} -s + \beta E[V(y|s; \mu)]; \quad (2)$$

where

$$E[V(y|s; \bar{x})] = \underbrace{F_Y(\bar{x}; s)}_{\text{Firm does not hire}} V(0) + \int_{\bar{x}}^z \underbrace{\frac{y}{1}}_{\text{Firm hires}} dF_Y(y; s); \quad (3)$$

and, by the definition of the bounded reservation quality,²⁹

$$\bar{x} = \max[(1 - \beta) V(0); z]; \quad (4)$$

Together, equations 2, 3 and 4 describe the model, and capture its key trade-offs. The firm has two distinct reasons to hire: (i) an extensive margin impact (by hiring, the firm avoids the loss of β), and (ii) an intensive margin impact (by hiring, the firm also gains y in every subsequent period). The firm invests in costly search activities ($c > 0$) in order to increase the number of matches λ and, therefore, to improve the expected quality of the top candidate.

Depending on the values of the key parameters, this model is capable of generating several different types of behaviour λ and, in particular, different comparative statics with respect to firm beliefs about worker quality.³⁰ In Figure 3, we consider a regime with particular relevance to our experimental results, using relatively large values for both β and z .³¹ On the horizontal axis of each panel, we show λ with higher values of β to the left, so that a move to the right indicates a more pessimistic firm belief.³² Panel A shows the firm's optimal choice of search effort, s . Panel B shows the firm's reservation quality, \bar{x} ; specifically, it shows both z , as a dotted blue line, and $(1 - \beta) V(0)$, as a dotted red line; the solid black line is therefore the upper envelope \bar{x} . Panel C shows the firm's anticipated probability of hiring, given its beliefs: $\Pr(y > \bar{x} | s)$. It is useful to visualise this anticipated probability in order to understand the motivation for the firm's optimal choice of both search effort and reservation quality.

Here, the model demonstrates three distinct regimes. First, on the left of each figure, β is relatively large; here, a 'pessimism shock' (that is, a decrease in β) leads to a decrease in search. This might be understood as the 'safe regime'; because the anticipated probability of finding a suitable applicant is extremely high (Panel C), the decrease in β causes a

²⁹ That is, if the firm were unconstrained, it would set \bar{x} such that $V(\bar{x}) = V(0)$. The firm chooses the greater of this value and the minimum reservation quality, z .

³⁰ We solve this model using a standard value function iteration (where we evaluate $E[V(y|s; \bar{x})]$ using a fast Monte Carlo integration). For our numerical implementation, we use the Exponential distribution for x ; that is, we use $F_x(x; \beta) = 1 - \exp(-\beta x)$.

³¹ For this illustration, we use $\beta = 1$, $\beta = 0.95$, $\beta = 50$ and $z = 5$.

³² Specifically, we graph values for β from $\beta = 3$ down to $\beta = 0.5$.

reduction in search in the current period (Panel A) because, when the pool quality is high and thus the expected value of next period is high, the marginal gain from searching more intensively today is small relative to waiting for a better applicant tomorrow. In this regime, the exogenous minimum quality (z) does not bind: the reservation quality is determined by the option value of leaving the position vacant to 'wait and see' whether a more suitable candidate can be found. In turn, this implies that the value function is continuous at \bar{x} (because $V(\bar{x}) = V(0) = \bar{x} = (1 - \beta)w$). In the safe regime, a small reduction in the expected pool quality reduces search.

At the other extreme, if β is very small (far right of the figure), the firm lies in an 'exit regime': the effort required to find a suitable candidate is so high that it is optimal not to search at all. In the exit regime, a pessimism shock has no effect at all since the firm has already decided not to search.

Between the two is an intermediate regime that we dub the 'unsafe regime'. The key characteristic of this regime is that β is not low enough as to make search unprofitable, but is sufficiently low that the exogenous minimum quality (z) binds on the firm's reservation quality (Panel B). This implies that the value function is discontinuous at \bar{x} : if the firm's top candidate is just below the appointable quality z , the firm suffers a discrete fall in profit compared to a candidate who barely reaches that threshold (formally, $\lim_{y \rightarrow \bar{x}^-} V(y) < V(\bar{x})$). In turn, this implies that the firm anticipates a meaningful probability that it will not hire (Panel C). In this regime, for sufficiently low β , the decrease in β causes an increase in search expenditure. Given the relatively high cost of not filling the position ($\beta \gg 0$), and driven by the firm's growing concern that it will not find a suitable candidate, the pessimism shock makes the firm search more intensely.

We view this intermediate regime as being not only the most interesting, but also empirically the most relevant.³³ The notion of an exogenous minimum candidate quality ($z \gg 0$) is justifiable by technical requirements (as in the earlier example, of a crane driver); it may also reflect organisational constraints, by which firms may face internal morale consequences of hiring underqualified candidates at a posted wage (see, for example, [Breza et al. \(2018\)](#)). The notion of a discrete cost of not filling a position ($\beta \gg 0$) resonates, for example, with 'O-Ring' style production processes ([Kremer, 1993](#)), in which the absence of a worker generates productivity costs for the firm as a whole.

What about jobseekers? The model discussed above describes a firm searching for a prospective employee. The same stylised setup can readily be understood in a jobseeker's perspective.

³³ The model is capable of generating other patterns, for other parameter values. For example, for $z = 0$, the firm lies in the 'safe regime' for all β .

mutandis, as describing search behaviour of our jobseekers. In the case of a jobseeker, \bar{x} and y can be understood as referring to posted wages, and $z > 0$ can be understood as capturing the disutility from being unemployed (including, for example, facing social pressure from family and friends, needing to impose on the generosity of others for financial and accommodation support, and so on). In the jobseeker context, we can think of $z > 0$ as a reference point below which the jobseeker is not willing to shift $\{$ driven, perhaps, by the leisure value of remaining unemployed, or by a distaste for low-status work (as documented, for example, by [Groh et al. \(2015\)](#) in Jordan).

To sum up, in the unsafe regime, the searching-firm model predicts that a pessimism shock (i) increases search effort, and (ii) has no effect on reservation quality $\{$ because, in that regime, $\bar{x} = z$. The model makes similar predictions for jobseekers: in the unsafe regime, they respond to a pessimism shock by increasing job search and, initially at least, they keep their reservation wage $\bar{x} = z$ unchanged. It is, however, not difficult to imagine a model extension in which z adjusts over time as unemployed workers 'swallow their pride' and settle for a less ideal wage $\{$ for example, as a result of liquidity constraints as the unemployed jobseeker exhausts available search funds and family support.

6 Evidence from belief elicitation surveys

Our theoretical framework shows $\{$ both on the firm side and the jobseeker side $\{$ that a sobering shock to beliefs can increase search intensity in a realistic setting. In the previous sections, we provided empirical evidence that workers and firms at the job fairs had over-optimistic beliefs about the quality of a likely match during the intervention and that search intensity by both firms and workers increased after the intervention ended. These findings are easily explained by our model.

To corroborate this interpretation, we returned to the field to collect new data on the beliefs of representative samples of firms and workers. Our objective is twofold. First, we want to provide independent evidence on the over-optimism of workers and firms in this setting. This will help dispel concerns that the experimental results are an artefact of either selection (the fairs may have exposed individuals to a negatively selected sample of the other side of the market) or misinterpretation (the invitation to the fair may have been misinterpreted as signal of ability/quality). Second, we want to collect data on a wider set of beliefs, in particular beliefs about the structural features of the labor market (e.g. the distribution of ability or wages). These beliefs are central in the definition of an optimal search strategy, but they are rarely investigated in surveys. Overall, we find clear

evidence that, without our intervention, beliefs are systematically overoptimistic on both sides of the market. This provides strong support for our interpretation of the channels by which the intervention impacts firms and workers. It also supports our hypotheses that unrealistic beliefs contribute to the low levels of hiring at the job fairs and that the downward revision of beliefs caused by the job fairs did bring expectations closer to reality.

6.1 The belief elicitation surveys

Following up on our initial experiment, in 2019 we conducted a new belief-elicitation exercise with firm managers and jobseekers. Since our objective was to understand whether potential misperceptions exist among market participants in the absence of our treatment, we did not go back to the original sample that took part in the experiment but rather surveyed a new representative sample of jobseekers and firms. Specifically, we contemporaneously sampled firms that were advertising vacancies on Addis Ababa's job boards and jobseekers that were looking for vacancies at those job boards.³⁴ The surveys have three unique features. First, they focus on a real, well-defined labour market. Second, they elicit beliefs on both sides of the market. While a number of papers study jobseeker beliefs, systematic data on the beliefs of firm managers is rare, especially in developing countries. Third, the surveys enable us to measure the accuracy of beliefs. In particular, we can contrast firms' answers with the true empirical counterparts obtained from the jobseeker survey and vice versa; this improves over existing studies that elicit beliefs but cannot measure their accuracy.³⁵

The questionnaire for firms carefully elicits firm managers' beliefs about the ability of jobseekers { a key element of our model. Since our intervention produced heterogeneous effects by jobseeker level of education, we document expectations with respect to tertiary-educated applicants and high-school graduates separately. Ability is measured as a jobseeker performance on a Raven's test. We took a number of steps to make sure

³⁴ We interviewed 395 firm managers and 779 jobseekers. We recruited jobseekers between the age of 18 and 29, who had at least a high school diploma. We contacted a random sample of firms that were advertising a position on the job boards or in the newspaper between the end of November and the end of December 2019. We also contacted some of the firms that jobseekers were applying to. In this way, we selected samples of firms and jobseekers that resemble on key dimensions our original experimental participants.

³⁵ In addition, our surveys focus on new samples of market participants as opposed to the original experimental subjects. This has two distinct advantages over working with our original samples. First, attrition after several years may have biased our samples. Second, the subjects that took part in the intervention were exposed to the information they gained at the fairs. Our interest here is in uncovering perceptions biases that exist in the absence of the fair intervention.

that lack of familiarity with such a test among firm managers does not distort our results. First, we provided the instructions for the test to the firm managers, so that they could familiarise themselves with it. Second, before managers answered the ability questions, we provided them with real statistics on the difference in test score between workers with a high (75th percentile) and an average GPA in our sample. This served the purpose of giving employers a sense of how test results correlate with an observable characteristic (GPA) commonly used in hiring, thereby providing an anchoring reference point. In addition to measuring expectations about jobseekers' ability, we also elicited managers' beliefs about the jobseekers' reservation wages and their work experience. This elicitation exercise was incentivised.³⁶

The questionnaire for jobseekers focuses on their reservation wage, their belief about the distribution of wages across sectors, and their belief about job-finding probabilities. We elicited beliefs about the distribution of wages by asking the jobseeker what proportion of jobs currently advertised pay a wage lower than a set of thresholds (from 10,000 ETB to 1,000 ETB per month). Similarly, we elicited reservation wages by asking the jobseeker whether they would accept a job that pays at least a certain amount. This amount was decreased until we found the wage bracket corresponding to the jobseeker's reservation wage. To minimise complexity, we did not incentivise the elicitation of beliefs among jobseekers. Finally, after the belief-elicitation was completed, jobseekers took a 12-item Raven test.

In light of our theoretical model, we use the data from the survey to investigate the hypothesis that firm managers and jobseekers may have distorted perceptions about the quality of available matches. For firm managers, we test this by comparing their perceptions about the quality of jobseekers (i.e., their ability) with the actual distribution of ability in our sample. For jobseekers, we test whether their beliefs about the expected duration of unemployment and the quality of available jobs (proxied by wages) align with reality. We also test how their reservation wages compare with prevailing wages.

6.2 Distorted beliefs among firm managers

From the belief elicitation survey of firms, we find clear evidence of distorted beliefs among managers. Our first result is that firm managers overestimate the ability of jobseekers. We asked firms to predict how many questions on a Raven's test can be answered correctly by a representative individual with high school or tertiary education, respectively. These

³⁶ Participants were told that one of the questions they were asked would be randomly drawn at the end of the survey and they would receive a prize based on the accuracy of their answer.

questions were asked after first familiarizing firm respondents with the test itself. Figure 4 shows that 65% of firms overestimate the average Raven's test scores for workers in the educational category for which they are currently hiring. Firms' average forecast of tests scores is higher than the true average – a result entirely driven by managers overestimating the ability of workers with tertiary education. This is illustrated in Figure B.2 through a series of 'raincloud plots' of managers' beliefs, with superimposed vertical lines showing the average jobseeker characteristic. Firm respondents also overestimate the ability gap between tertiary and secondary-educated jobseekers. The average secondary-educated and tertiary-educated jobseekers correctly answer 5 and 5.3 questions on the Raven test, respectively. By contrast, the median firm forecasts that secondary-educated jobseekers correctly answer 4 questions vs. 6 questions for tertiary-educated jobseekers. In other words, the true ability premium associated with tertiary education is less than one fourth of what firms expect. Furthermore, the perceived difference between the two groups is twice as large as the difference in Raven performance between individuals at the 75th percentile and at the mean of the GPA distribution – the anchoring information we gave to firms before these forecasts. Overall, almost 75% of firms overestimate the ability of tertiary-educated jobseekers and about 90% of them overestimate the ability premium associated with tertiary education. Because most firms in our sample are trying to hire workers with tertiary education, the average firm in the sample overestimates the ability of the types of worker they are trying to hire – a finding that is summarised in Figure 4.

Additionally, we find that firms overestimate work experience and reservation wages among tertiary educated jobseekers. More than 75% of managers overestimate the share of tertiary-educated jobseekers who have at least two years of work experience (Figure B.2). The median manager expects about 20 percent of tertiary-educated workers to have two years of work experience, a figure that is almost twice the actual proportion. Furthermore, in the survey, we ask firm respondents to indicate the proportion of jobseekers who would accept different wage levels for the most common job available at the firm. Figure 5 shows that firms overestimate the reservation wages of tertiary-educated jobseekers across occupations – but most starkly with respect to professional roles. Finally, as was the case for ability, the patterns for secondary-educated workers are reversed: firms underestimate both their work experience (e.g. Panel B of Figure B.2) and their reservation wages (Figure B.3).

In sum, firms on average overestimate the ability and work experience of the jobseekers they are trying to hire. This shows unequivocally that beliefs are inaccurate in a representative sample of recruiters. These findings thus help explain why hiring at the

job fairs was modest: firms were disappointed by the experience level of the tertiary-educated jobseekers they met and hence made few if any offers; at the same time, they underestimated the ability and reservation wages of the secondary-jobseekers they met and, consequently, made offers to this group that were rejected. This interpretation is further supported by clear evidence that past experience is highly sought after by firms. The most common reason firms report for not hiring more at the fair is 'insufficient work experience' (34% of firms).³⁷ Furthermore, dyadic analysis of firm-requested meetings shows that past experience is the strongest predictor of firms' meeting requests (and the effect is strongest among workers with tertiary education, in line with the rest of the evidence).³⁸

Overall, the evidence from the belief-elicitation firm survey is thus consistent with the hypothesis that the firms that came to the job fairs had incorrect beliefs about the experience levels of tertiary-education candidates and about the quality and reservation wage of secondary-educated candidates. It is therefore reasonable to assume that, by allowing firms to interact with a large number of such candidates at once, the fairs gave them an opportunity to update their beliefs. This is indeed what we showed earlier in this paper: Section 4.2.3 documented that treated firms became more pessimistic about the number of workers they would hire in the next 12 months; and Table 6 showed that treated firms become more pessimistic about hiring both types of workers, but especially those in occupations that typically require tertiary education. This, in turn, led firms to revise their search strategy in the direction predicted by our model, that is, by increasing their search effort.

6.3 Distorted beliefs among jobseekers

Turning to jobseekers, the belief elicitation survey provides clear evidence of over-optimism on the probability of finding a good job. The evidence is particularly strong for secondary-educated jobseekers. These jobseekers overestimate both the probability of obtaining a permanent job with an open-ended contract, as well as the probability of obtaining a professional job. We discuss each of these findings in turn.

³⁷ Other common reasons relate to the perceived expertise of workers or poor interview performance (see Table B.18). Educational mismatch plays a role, but is certainly not the most important factor.

³⁸ We apply the same dyadic regression approach as in equations (5) and (6) and report the results in Table B.19. The dependent variable is $isrequest_{fj}$, a dummy equal to one if firm f requested a meeting with jobseeker j , using a centralized meeting-request algorithm that we offered to firms at the fair. Regressors include jobseeker and firm characteristics. The results are not driven by firms who sought experienced jobseekers outside the fair: even firms willing to hire graduates without work experience at baseline are more likely to request experienced jobseekers at the fair.

First, we present evidence on the probability of securing a permanent job. Figure 7 compares jobseekers' perceived likelihood of finding a permanent position with data on the actual likelihood of getting a permanent position in the control group of our experimental sample.³⁹ We see that 55% of jobseekers with a high-school diploma expect to find a job with a permanent contract in less than 1 year. In reality, only 5.8% of such candidates on our experimental sample found a permanent job within 1 year. Furthermore, when asked about jobseekers in the same age cohort with the same education and work experience, respondents expect 30% of them to find a permanent job within one year. This suggests that jobseekers with a high-school diploma are not only over-confident about their own ability relative to the rest of their cohort, they are also over-optimistic about the average prospects of individuals like themselves.

Second, we present evidence on the probability of securing a professional job. In the survey, we asked jobseekers who were targeting high-skilled jobs in managerial, technical or professional positions, who they thought would eventually get that job. Among jobseekers who did not complete a tertiary degree, less than 20% believe that the vacancy would eventually be filled by someone with a tertiary degree. Furthermore, only a fifth of them think that a tertiary degree is a requirement for the job. This differs markedly from that is reported by the firms participating in the belief elicitation survey: 71% of high-skilled vacancies have degrees as a minimum requirement, and 74% eventually go to someone with a tertiary degree. Similarly, half of the jobseekers think that no previous work experience is required for securing a high-skilled job and 36% think that the job will go to someone without any formal work experience. In reality, surveyed firms report that only 16% of high-skilled jobs require no formal experience while 59% of these jobs require two or more years of experience.

On the other hand, jobseekers' beliefs about the distribution of wages paid for available jobs are well-aligned with reality. When asked about average wages in different occupations, the answers respondents provide closely track prevailing wages in different sectors (Figures B.4 and B.5). Jobseekers with secondary education also have fairly accurate beliefs about the distribution of wages in specific occupations (Figure B.4). The

³⁹ Since we only interview jobseekers once, we do not have data on the length of their unemployment spell in the most recent survey and we have to resort to the experimental sample. We believe this provides a valid benchmark. First, the two populations were selected using similar screening criteria based on age and education. Thus, when we re-weight by observables to ensure comparability between the two samples, there are no qualitative changes in our findings. Second, although the two samples were interviewed a few years apart, aggregate labour market conditions are not significantly different between the two periods. Third, to drive the observed divide between more educated jobseekers and less educated ones, labour market conditions should have varied differentially for different groups of workers. We have no evidence of that occurring.

same is true for jobseekers with tertiary education who, if anything, overestimate the proportion of jobs at the bottom of the wage distribution (Figure B.5).

What are the implications of these beliefs for the search strategies chosen by jobseekers? We find that jobseekers, especially those with secondary education, set unrealistic reservation wages and target jobs that they are unlikely to get. In terms of reservation wages, we find that 70% percent of jobseekers with only a high-school diploma and no permanent work experience state they would reject a job paying 2000 ETB per month, even though 44% of jobs for that occupation and level of experience pay less than 2000 ETB per month. Jobseekers with secondary education also often seek positions for which firms largely hire tertiary educated jobseekers and which, therefore, they are unlikely to get. Figure 6 shows this clearly. A large proportion of jobseekers with secondary education (Panel A) seek employment in professional categories such as 'Technicians and professionals' and 'Services and sales workers', even though firms offer few opportunities in those roles to jobseekers with only secondary education. Overall, these findings are consistent with the behavior of secondary-educated jobseekers at the job fairs: these jobseekers rejected job offers that were largely for non-professional jobs and paid less than their high reservation wages.

Why young jobseekers hold incorrect beliefs is unclear, but one possible explanation is that their expectations partly reflect the experiences of older relatives who entered the labor market at a time when tertiary education was more scarce and hence a larger share of positions was open to those with secondary education.

6.4 Alternative explanations

The belief elicitation surveys have shown that both firms and jobseekers hold unrealistic beliefs about each other. We have argued that this provides a plausible explanation for the changes in search strategies induced by the job fairs. It can also account for the improved employment outcomes at endline despite the fact that the job fairs did not, by themselves, match jobseekers to jobs. We now examine two alternative explanations and check whether they find some support in our data.

One alternative explanation is selection effects, namely, that firms and jobseekers had correct beliefs but the fairs exposed them to non-representative samples of market participants. As a result, workers and firms misinterpreted their experience at the fairs as a negative signal about the state of the labour market. For instance, the firms that attended the fairs may have had more competitive vacancies than the average firm on

the market. If workers did not take this selection into account, they may incorrectly have become more pessimistic about the probability of securing a job and this, our model predicts, induced them to search harder. Similarly, the jobseekers attending the fairs may have been negatively selected. If the firms did not realize this, they would have become less optimistic about finding workers matching their experience requirements or wage offers, thus leading them to search more intensively.

Another alternative explanation is that matching at the fairs was of low quality, possibly due to poor logistics or to difficulties for workers to locate employers, and this low matching quality was erroneously interpreted by participants as a signal of poor labour market fundamentals. This, in turn, would have induced firms and/or workers to search harder as a result of being invited or attending the fairs.

These alternative explanations share one feature in common: the job fairs generated a misleading signal about the market that moved participants' initially correct beliefs away from the truth. To reject these explanations, we first note that the belief elicitation surveys demonstrate that incorrect beliefs are held by both firms and jobseekers unexposed to the job fairs. Still, it could be that the job fairs moved these incorrect beliefs even further away from the reality of the market because of selection effects or poor matching. Descriptive evidence on selection and matching at the job fairs suggests that this was not the case. Earlier in the paper we showed that the samples invited to the fairs were fairly representative and that take-up was not highly selected. This addresses the selection concern. Regarding matching, we presented in Section [A.1](#) and Section [A.2](#) results showing that the meetings that took place at the fairs were positively selected on expected match quality. This rules out explanations based on poor matching. Instead, the evidence points to the conclusion that attendees came to the job fairs with over-optimistic beliefs similar to those we document in our belief elicitation surveys, and that the fair had a sobering effect that corrected these misperceptions. Based on this, we conclude that the combined evidence we gathered in the experiment and the belief elicitation surveys is most consistent with our preferred interpretation based on our theoretical model, namely: both sides of the markets acquired new information about the fundamentals of the labor markets through the rich interactions generated by the fairs, this caused them to adjust to their beliefs to be more in line with reality, and this in turn led them to change their search strategy and, for jobseekers, get better employment outcomes.

7 Conclusion

We run a novel experimental job fair, with a unique dual randomization { both on the side of jobseekers and of participating firms. The invited jobseekers are representative of the young jobseekers whom firms usually hire, and participating firms are a representative sample of large employers. We facilitate interactions between jobseekers and firms by providing information about jobseekers' education and firms' vacancies, and by suggesting matches based on a Gale-Shapley algorithm. We study both the immediate effects of the treatment on jobseekers' and firms' outcomes, and subsequent effects on both search strategies and expectations.

We find that the fairs generate a rich set of interactions between jobseekers and firms, and that the matching algorithm is successful in increasing the efficiency of the matching process. The immediate impact of the treatment on employment outcomes is limited with few hires made at the fair, but we find clear evidence of delayed effects of treatment, as both firms and jobseekers learn from the information they acquire at the fairs: they change their expectations accordingly and adjust their search strategies. Treated jobseekers with at most a high-school diploma had misaligned reservation wages prior to treatment; after the fairs, they search harder, lower their reservation wages, and experience a significant increase in their probability of obtaining a formal job at endline (6 months after the second fair). Treated firms increase their search efforts and substitute away from tertiary educated workers on whom they had overly optimistic expectations at baseline. A follow-up belief-elicitation exercise with similar jobseekers and firms confirms both that firms have overly optimistic perceptions of the skills of jobseekers with tertiary education, and that jobseekers have overly optimistic beliefs about the probability of obtaining professional jobs given their qualifications.

The main contribution of our paper is to show that both firms and jobseekers hold inaccurate beliefs about market fundamentals { that is, labour market participants suffer not merely from a problem of information asymmetries, but from deeper misperceptions of the distribution of important traits among other market participants. We find that by facilitating rich interactions between the two sides of the market the fairs serve to reduce these deep misperceptions. Our results show that active labour market policies that increase contact between jobseekers and firms { such as job fairs, and including many other classes of policy intervention { are likely to generate important learning effects on both sides of the market even when immediate impacts on employment outcomes are limited.

References

- Abebe, G., A. S. Caria, M. Fafchamps, P. Falco, S. Franklin, and S. Quinn (2021). Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City. *The Review of Economic Studies* 88(3), 1279{1310.
- Abebe, G., A. S. Caria, and E. Ortiz-Ospina (2021). The Selection of Talent: Experimental and Structural Evidence from Ethiopia. *American Economic Review* 11(6), 1757{1806.
- Alfonsi, L., O. Bandiera, V. Bassi, R. Burgess, I. Rasul, M. Sulaiman, and A. Vitali (2020). Tackling Youth Unemployment: Evidence from a Labor Market Experiment in Uganda. *Econometrica* 88(6), 2369{2414.
- Bandiera, O., V. Bassi, R. Burgess, I. Rasul, M. Sulaiman, and A. Vitali (2021). The Search for Good Jobs: Evidence from a Six-year Field Experiment in Uganda. Working paper.
- Banerjee, A. and S. Sequeira (2021). Learning by Searching: Spatial Mismatches and Imperfect Information in Southern Labor Markets. Working paper.
- Bassi, V. and A. Nansamba (2021). Screening and Signaling Non-cognitive Skills: Experimental Evidence from Uganda. *Economic Journal*.
- Benjamini, Y., A. M. Krieger, and D. Yekutieli (2006). Adaptive Linear Step-up Procedures that Control the False Discovery Rate. *Biometrika* 93(3), 491{507.
- Bobotas, P. and M. V. Koutras (2019). Distributions of the minimum and the maximum of a random number of random variables. *Statistics & Probability Letters* 146 57{64.
- Breza, E., S. Kaur, and Y. Shamdasani (2018). The morale effects of pay inequality. *The Quarterly Journal of Economics* 133(2), 611{663.
- Bruhn, M. and D. McKenzie (2009). In Pursuit of Balance: Randomization in Practice in Development Field Experiments. *American Economic Journal: Applied Economics* 1(4), 200{232.
- Bryan, G. and M. Morten (2019). The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia. *Journal of Political Economy* 127(5), 2229{2268.
- Cameron, A. C., J. B. Gelbach, and D. L. Miller (2011). Robust Inference with Multiway Clustering. *Journal of Business & Economic Statistics* 29(2), 238{249.

- Caria, S. and P. Falco (2021). Sceptical Employers: Experimental Evidence on Biased Beliefs Constraining Firm Growth. Working paper.
- Carranza, E., R. Garlick, K. Orkin, and N. Rankin (2021). Job Search and Hiring with Limited Information about Workseekers' Skills. Working paper.
- DellaVigna, S., A. Lindner, B. Reizer, and J. F. Schmieder (2017). Reference-dependent job search: Evidence from Hungary. *The Quarterly Journal of Economics* 132(4), 1969{2018.
- Eeckhout, J. (2018). Sorting in the Labor Market. *Annual Review of Economics* 10(1), 29.
- Franklin, S. (2018). Location, Search Costs and Youth Unemployment: Experimental Evidence from Transport Subsidies. *The Economic Journal* 128(614), 2353{2379.
- Gale, D. and L. S. Shapley (1962). College Admissions and the Stability of Marriage. *The American Mathematical Monthly* 69(1), 9{15.
- Groh, M., D. McKenzie, N. Shammout, and T. Vishwanath (2015). Testing the Importance of Search Frictions and Matching Through a Randomized Experiment in Jordan. *IZA Journal of Labor Economics* 4(1), 1{20.
- Hensel, L., T. Tekleselassie, and M. Witte (2022). Formalized employee search and labor demand. Technical report, Working Paper.
- Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019). The Allocation of Talent and US Economic Growth. *Econometrica* 87(5), 1439{1474.
- Kremer, M. (1993). The O-ring Theory of Economic Development. *The Quarterly Journal of Economics* 108(3), 551{575.
- Krueger, A. B. and A. I. Mueller (2016). A Contribution to the Empirics of Reservation Wages. *American Economic Journal: Economic Policy* 8(1), 142{179.
- Lee, D. S. (2009). Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects. *The Review of Economic Studies* 76(6), 1071{1102.
- McCall, J. J. (1970). Economics of Information and Job Search. *The Quarterly Journal of Economics* 113(1), 126.
- Mueller, A. I., J. Spinnewijn, and G. Topa (2021). Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence and Bias. *American Economic Review* 111(1), 324{63.

- Rogerson, R., R. Shimer, and R. Wright (2005). Search-theoretic Models of the Labor Market: A Survey. *Journal of Economic Literature* 43(4), 959{988.
- Singh, N., N. Fernando, and G. Tourek (2022). Hiring frictions and the promise of online job portals: Evidence from india. Technical report, Working Paper.
- Spinnewijn, J. (2015). Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs. *Journal of the European Economic Association* 1(3), 130{167.
- Terviö, M. (2009). Superstars and mediocrities: Market failure in the discovery of talent. *The Review of Economic Studies* 7(2), 829{850.
- Van den Berg, G. J. and B. van der Klaauw (2019). Structural empirical evaluation of job search monitoring. *International Economic Review* 60(2), 879{903.
- van der Klaauw, B. and L. Ziegler (2021). A Field Experiment on Labor Market Speed-dates for Unemployed Workers. Technical report, Working paper.
- Wilken, W. (2021). Job search with spatial and informational frictions: Structural estimation on Ethiopian data. Master's thesis, University of Oxford.
- Wright, R., P. Kircher, B. Julien, and V. Guerrieri (2021). Directed Search and Competitive Search: A Guided Tour. *Journal of Economic Literature*, 90{148.

Tables

Table 1: Firm size by sector

Industry	Client services	Worker Type			White collar	All workers	Sample Size
		Production	Support sta	Support sta			
Construction, Mining, Farming	2.7	92.7	21.7	21.8	143.2	92	
Tours-Hospitality	15.8	7.4	13.2	7.4	46.4	102	
Finance, Services, Retail	146.6	33.7	96.6	183.3	473.3	104	
Education, Health, Aid	12.6	5.2	31.2	73.6	131.0	126	
Manufacturing	24.4	149.0	37.4	33.7	250.2	69	
All Industries	26.9	52.4	33.1	52.8	171.5	493	

Notes: This table describes the firms in our sample, disaggregating by primary sector and by type of occupations.

Table 2: Mismatched expectations: Reservation wages of workers before the fairs, wages offered at the fairs, and endline employment outcomes (Medians)

	Education of worker			
	High-school	Vocational	Diploma	Degree
Panel A: Workers' reservation wages, and firms' wages for jobs offered at the job fairs				
Worker reservation wages before fairs				
With experience (13%)	1500	2000	2000	3000
Without experience (87%)	1300	1500	1600	2500
Firm wages for positions at fairs				
Require experience	1588	1900	3250	5685
Don't require experience	855	1018	1168	3500
All jobs	973	1500	2900	4500
Panel B: Workers' expectations at baseline				
Expects at least one job offer in next four months				
All experience levels	73%	72%	74%	72%
With experience	81%	72%	82%	72%
Without experience	73%	72%	72%	73%
Expected wage "for someone like me one" year from now				
All experience levels	1400	1600	1800	2500
With experience	1500	1900	1900	3400
Without experience	1400	1500	1800	2500
Panel C: Workers' employment outcomes at endline (control group)				
Worker employment rates at endline				
All jobs	50%	56%	43%	69%
Permanent jobs	6%	17%	19%	35%
Worker wages at endline by experience				
With experience	1450	1450	1743	3000
Without experience	975	1400	1350	2100
All experience levels	1000	1400	1500	2300

Notes: This table describes self-reported reservation wages (for jobseekers) using phone survey data in the weeks just prior to the first job fair, offered wages at the job fair (for firms), and endline wages (for jobseekers in the control group), disaggregated by types of worker and type of job.

Table 3: Firm recruitment methods

Outcome	Estimated ITT	Control Mean	Observations
Firm performed formal interviews (professionals)	0.0440 (.038) [.138]	0.682	473
Firm performed formal interviews (non-professionals)	-0.0140 (.039) [.401]	0.607	473
Did any advertising for new hires	0.0580 (.032)* [.074]*	0.789	473
Did advertising for professional positions	0.120 (.038)*** [.009]***	0.595	473
Did advertising on the job boards	0.0960 (.042)** [.044]**	0.331	473

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q -values are reported in square brackets.

Table 4: Workers' job search and employment outcomes after the fairs

	(1)	(2)	(3)	(4)	(5)
	Board visits	Worked	Per. work	Formal job	Earnings
Panel A: Average Treatment Effect					
Fairs	3.012** (1.285)	-0.004 (0.0273)	0.024 (0.0183)	0.026 (0.0193)	34.110 (75.89)
Observations	1,705	1,705	1,705	1,705	1,690
R-squared	0.006	0.000	0.001	0.001	0.000
Control Mean	14.780	0.562	0.171	0.224	0.224
Panel B: Treatment Effect by Education					
Fairs*HighSchool	4.197** (1.719)	-0.012 (0.0397)	0.0576** (0.0262)	0.0482* (0.0275)	0.421 (86.82)
Fairs*PostSecondary	1.251 (1.335)	0.006 (0.0313)	-0.026 (0.0234)	-0.008 (0.0236)	84.360 (128.6)
Observations	1,705	1,705	1,705	1,705	1,690
R-squared	0.008	0.000	0.005	0.003	0.000
ControlMean HighSchool	10.980	0.508	0.058	0.108	0.108
ControlMean PostSecond	16.550	0.587	0.223	0.277	0.277
Test High=Post (p)	0.118	0.718	0.018	0.116	0.577

Notes: Each row reports a separate regression. 'Wage mismatch' refers to the absolute difference between the worker's reservation wage (in logs) and the expected wage for a worker of that skill/education level (in logs). For each regression, we report the estimated ITT from participating in the job fair (disaggregated, in Panel B, by whether the worker has post-secondary education or merely high school). Standard errors are reported in parentheses. In the bottom row, we report p-values for a test of the null hypothesis that the effect of treatment is equal between high-school and post-secondary sub-samples.

Table 5: Workers' expectations after the fairs

	(1)	(2)	(3)	(4)	(5)	(6)
	Res Wage	Wage Mismatch	High wage	O ers expected	Want perm job	Likelihood job
Panel A: Average Treatment Effect						
Fairs	-0.0660* (0.0369)	-0.0422* (0.0224)	-0.025 (0.0436)	-0.090 (0.139)	-0.0366*** (0.0132)	-0.015 (0.0689)
Observations	1,501	1,501	1,516	1,680	1,702	1,543
R-squared	0.005	0.003	0.001	0.001	0.008	0.000
Control Mean	7.417	0.529	8.206	1.421	0.968	3.013
Panel B: Treatment Effect by Education						
Fairs*HighSchool	-0.0867* (0.0486)	-0.0714** (0.0337)	-0.047 (0.0567)	-0.165 (0.156)	-0.0580*** (0.0190)	-0.040 (0.0881)
Fairs*PostSecondary	-0.035 (0.0352)	0.001 (0.0219)	0.007 (0.0381)	0.022 (0.165)	-0.005 (0.0122)	0.024 (0.0755)
Observations	1,501	1,501	1,516	1,680	1,702	1,543
R-squared	0.005	0.006	0.001	0.002	0.011	0.000
ControlMean HighSchool	7.183	0.561	7.996	1.332	0.983	3.009
ControlMean PostSecond	7.522	0.514	8.300	1.462	0.961	3.015
Test High=Post (p)	0.278	0.063	0.283	0.240	0.010	0.506

Notes: Each row reports a separate regression. `Wage mismatch' refers to the absolute difference between the worker's reservation wage (in logs) and the expected wage for a worker of that skill/education level (in logs). For each regression, we report the estimated ITT from participating in the job fair (disaggregated, in Panel B, by whether the worker has post-secondary education or merely high school). Standard errors are reported in parentheses. In the bottom row, we report p-values for a test of the null hypothesis that the effect of treatment is equal between high-school and post-secondary sub-samples.

Table 6: Firms' expected number of future hires after the fairs

Outcome	Job Fair	Control Mean	N
Panel A: Aggregate hiring expectations			
All	-13.33 (7.049)*	45.63	472
Panel B: Hiring expectations by occupational type			
Higher-education positions	-9.691 (5.337)*	33.87	419
Lower-education positions	-5.051 (3.55)	14.66	415

Notes: The outcome variable is the number of workers the firm expects to hire in the following 12 months in different occupations. Higher Educ. includes "Professional/Managerial" and "Client Service" occupations, which typically require higher levels of education. Lower Educ. includes "Production Workers" and "Support Service" occupations. All includes workers in all occupational categories. Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses.

Figures

Figure 1: Output of the matching algorithm

Notes: This figure illustrates the outcome of the matching algorithm. Each point represents a stable match recommended by the algorithm. The figure shows which combinations of firm rankings and job-seeker rankings generated these recommended matches. The graph provides a visual illustration that the algorithm worked well in the sense of generating matches between firms and job-seekers who are, on the basis of job-seeker skills and experience, reasonably well-suited to each other.

Figure 2: Impacts on Job Search by Fortnight

Notes: This figure shows the probability, for each fortnight, that treated job-seekers visit the job board, relative to job-seekers in the control group. Fortnight 0 is when the first job fair was held; the second fair was held in fortnight 8. We estimate the difference in probabilities using a Linear Probability Model in an ANCOVA specification, in which we regress job search on treatment, baseline search status and a vector of baseline balancing variables. We cluster at the level of individual job-seekers, and show both point estimates and 90% confidence intervals; we do this both by regressing on fortnight dummies, and by imposing a quadratic shape.

Figure 3: Model predictions

panel a: search

panel b: reservation quality

panel c: anticipated probability of hiring

This figure shows the key predictions of our dynamic search model. Specifically, it shows numerical results obtained by a value function iteration (and using a standard Monte Carlo integration). We use the Exponential distribution: $F_x(x; \lambda) = 1 - \exp(-\lambda x)$, and set $\lambda = 1$, $\beta = 0.95$, $\gamma = 50$ and $z = 5$. The horizontal axis shows values for θ from $\theta = 3$ down to $\theta = 0.5$. The panels respectively show (i) the firm's optimal search effort, s (Panel A), (ii) its reservation quality, \bar{x} (Panel B), and (iii) the resulting probability of hiring, $\Pr(y \geq \bar{x} | s)$ (Panel C). (Panel B shows both \bar{x} , as a dotted blue line, and $(1 - \beta) V(0)$, as a dotted red line; the solid black line is therefore the upper envelope.)

Figure 4: Distribution of manager's forecast error on jobseekers' average Raven's test score

Notes: The forecast error is computed as the difference in percentage points between a manager's belief about the average score of workers in the educational category (high school or tertiary education) most sought after among current vacancies open at the firm, and the actual average score of workers in that educational category.

Figure 5: Firms' beliefs about the distribution of job-seekers' reservation wages: Tertiary-educated job-seekers

Notes: These figures show the distribution of r_m beliefs about job-seekers' reservation wages. In each graph, we show this at five wage points: in each case, we show the distribution of r_m beliefs (using thin bars to show the 10th and 90th percentiles, thick bars to show the 25th and 75th percentiles, and a star symbol to show the median). The coloured lines show the true proportion of our job-seekers with a given reservation wage (where the shaded area represents the 90% confidence interval for the proportion).

Figure 6: Distribution of occupation sectors
panel a: high-school graduates

panel b: tertiary graduates

Note: This figure shows the distribution of (i) the proportion of total jobs in the most common occupations in each firm, and (ii) the sector of the job most commonly looked at by job-seekers in the last week. We show both bars for the five most common sectors for the firm side. We report D_{KL} , the Kullback-Leibler distance from the distribution of jobseeker sectors to the distribution of firm occupation sectors.

Figure 7: Jobseekers' expectations of finding a job with a permanent contract in the next 12 months

Note: `Expectation: Self' refers to jobseekers' stated probabilities that they will be employed with a permanent contract in the next 12 months, as measured in our 2019 follow-up survey. `Expectation: Other' refers to jobseekers' stated probabilities that others like them will be employed with a permanent contract in the next 12 months, as measured in our 2019 follow-up survey. `Actual' refers to the actual proportion of jobseekers who found a job with a permanent contract, using our original survey data.

Online Appendix

A Robustness Checks

A.1 Was the market at the job fairs too thin?

One possible explanation for this small direct effect is the market at the job fairs was too thin: there were too few high-quality worker-firm matches available. We present evidence against this hypothesis both from the jobseeker and the firm side. First, we investigate whether the jobs on offer were too few or did not match jobseekers' interests. To study this issue, we use data that was collected from participating firms prior to arriving at the fairs. Firms were to provide a roster of all their open vacancies at the time of the fairs.⁴⁰ The average firm at the fair had two vacancies open and was looking to hire seven workers. 70% of participating firms had at least one vacancy. In total, there were 711 vacancies and 1,751 jobs available at the fairs. The occupational composition of the vacancies exhibits considerable overlap with the distribution of occupations desired by jobseekers invited to the job fairs. It is therefore unlikely that firms did not have enough vacancies of the kind that jobseekers wanted.

Second, we investigate whether jobseekers were negatively selected and hence firms were reluctant to hire them. To explore this possibility, we compare the jobseekers who attended (about 60% of those invited) to those in the full sample, which is near-representative of educated young jobseekers in Addis Ababa at the time of the study. In Appendix Table B.6 we regress attendance at the fairs on a rich set of baseline characteristics. We find no evidence suggesting that observably weaker candidates are more likely to attend the fairs: education and current employment do not significantly predict attendance. The only two robust predictors of attendance are instead associated with a positive motivation to work: attendance is higher among those jobseekers who search

⁴⁰ We define a vacancy as an open position for a specific occupation. Firms first produced a list of vacancies (e.g. a firm could report that they were both looking for clerical workers and for drivers) and, then, for every vacancy, they reported the number of workers they were planning to hire in that position.

the most at baseline and who produce a formal certificate to employers⁴¹. Further, in the second job fair, we showed firms the list of qualifications of jobseekers at the fair and asked them whether they were interested in interviewing some of them. Most responded positively and provided the names of several candidates of interest to them. Across both fairs, firms report meeting 20 jobseekers on average. We can therefore rule out that firms were in principle uninterested in the jobseekers that attended the fairs.

A.2 Did the fairs suffer from congestion and mis-coordination?

Since both employers and jobseekers were interested in each other and willing to interact, could the small direct effect of the fairs be due to congestion and miscoordination? That is, could the effect be explained by firms and jobseekers having wasted their time and effort talking to the wrong people? To investigate this possibility, we test whether the jobseeker-firm pairs that met are those that were most suitable for each other, given the mix of employers and jobseekers at the fairs. We use two types of variables to assess mutual suitability: the synthetic rankings, and the proposed matches that we suggested to participants. The two ranking variables are $Rank_{kj}$, which is firm f 's ranking of jobseeker j , and $Rank_{kf}$, which is jobseeker j 's ranking of firm f . The two proposed match variables are: $Gale_Shapley_{fj}$, which is equal to 1 if jobseeker j and firm f were recommended to each other by our Gale-Shapley algorithm, and $Random_{fj}$, which is 1 if jobseeker j and firm f were randomly recommended to each other by us. If firms and jobseekers are able to engage in promising interactions, we expect participants' rankings to predict who wishes to meet with whom and who actually meets whom. If our matching algorithm was capable of identifying promising matches instead of random matches, we expect meetings and willingness to meet to be predicted by $Gale_Shapley_{fj}$ but not by $Random_{fj}$.

To test these hypotheses, we estimate two dyadic regression models:

$$y_{fj} = \alpha_0 + \alpha_1 Rank_{kj} + \alpha_2 Rank_{kf} + \epsilon_{fj}; \quad (5)$$

$$y_{fj} = \alpha_0 + \alpha_1 Gale_Shapley_{fj} + \alpha_2 Random_{fj} + \epsilon_{fj}; \quad (6)$$

where y_{fj} is either $request_{fj}$, a dummy equal to one if firm f requested a meeting with jobseeker j , or $meet_{fj}$, which equals one if firm f and jobseeker j actually met. Standard errors are clustered two-way at the level of the firm and jobseeker (Cameron, Gelbach, and Miller, 2011).

We report estimates in Table A.1, using the jobseekers and firms who attended the fairs. We find that the synthetic rankings predict both requested meetings and actual meetings. The effects are large and significant. Moving from the highest to the lowest rank is associated with an almost 100 percent decrease in the probability of a requested meeting, and about a halving of the probability of an actual meeting. We interpret these

⁴¹ Invitees already in permanent employment at the time of the fairs are slightly less likely to attend. But the effect is unlikely drive our results: 4% of those attending the fairs have a permanent job compared to only 5.6% of the total sample.

results as showing that the fairs are effective in bringing together jobseeker-firm pairs who { at least on the basis of observable characteristics } value each other. Algorithmic recommendations are also shown to have a strong predictive power: matches suggested by our algorithm are about 200 percent more likely take place than non-suggested matches. In contrast, the coefficient on randomly suggested matches is small and never significant. This contrast suggests that our stylized matching algorithm was useful in identifying matches that were deemed worth pursuing by market participants. The fairs thus appear to have reached their objective of facilitating meetings between jobseekers and the firms that suited them best.

This interpretation is supported by comparing the application-to-interview and interview-to-offer rates at the fairs versus in the open market. First, in the open market, job-seekers secured an interview for every 3.5 job applications, an offer for every 1.9 interviews, and a job for every 3.3 interviews over the period between the baseline and endline surveys. This implies that contacts with employers at the fair (20 on average) were much less likely to result in an interview than a formal job application. The contrast is particularly striking for highly educated job-seekers, who tend to do better in the labour market but did particularly poorly at the fair. Second, the 1.4 conversion rate of interviews into offers compares favorably to the 1.9 conversion rate observed outside the fairs. Third, the conversion of interviews into jobs is much lower at the fair: one job for 7.5 interviews instead of 3.5 outside the fairs. A large majority (81%) of offers made in the aftermath of the fairs were rejected. To verify these findings, we conducted a phone survey of firms immediately after each job fair. Appendix Tables B.12 and B.13 show the immediate impact on overall hiring and the type of job candidate hired, respectively. These results confirm that the fairs had no significant impact on short-term hiring by treated firms.

Appendix Table A.1: Dyadic regressions: Rankings, matches and meetings

	Requested (1)	Actual (2)	Requested (3)	Actual (4)	Requested (5)	Actual (6)
Firm ranking of workers	-0.006 (.001)	-0.002 (.0006)			-0.006 (.001)	-0.001 (.0006)
Worker ranking of rms	-0.002 (.002)	-0.001 (.002)			-0.002 (.002)	-0.001 (.002)
Algorithm suggestion			.020 (.007)	.015 (.006)	.014 (.006)	.014 (.006)
Random suggestion			.0006 (.006)	.003 (.007)	.0009 (.006)	.003 (.007)
Const.	.027 (.004)	.012 (.004)	.012 (.001)	.006 (.001)	.026 (.004)	.011 (.003)
Obs.	27778	27778	27778	27778	27778	27778
Effect size: max to min rank	.024	.006			.024	.005
Algorithm = Random			.029	.14	.123	.178

Notes: This table reports the estimates of equations 5 and 6. The highest ranked worker and firm are assigned a value of zero. Lower ranks correspond to higher numbers. Standard errors are corrected for two-way clustering at the level of the worker and at the level of the firm. The last row reports the p-value of an F-test of the hypothesis that the effect of the algorithmic and the random suggestion are the same.

B Additional Figures and Tables

Appendix Table B.1: Summary at baseline and tests of balance

	(1) Control Mean	(2) (SD)	(3) Job Fairs	(4) N	(5) F-test P
Degree	0.18	0.39	-0.01 (0.62)	1829	0.619
Vocational	0.43	0.49	-0.00 (0.91)	1829	0.910
Employed	0.31	0.46	-0.04 (0.15)	1829	0.155
Searched for work	0.50	0.50	-0.01 (0.76)	1829	0.763
Diploma or degree	0.25	0.43	-0.00 (0.99)	1829	0.993
Female	0.52	0.50	0.01 (0.85)	1829	0.848
Born outside of Addis Ababa	0.37	0.48	-0.03 (0.46)	1829	0.459
Amhara ethnic group	0.46	0.50	-0.02 (0.59)	1829	0.590
Oromo ethnic group	0.26	0.44	-0.04 (0.17)	1829	0.171
Worked in the last 6 months	0.46	0.50	-0.04 (0.19)	1829	0.186
Married	0.20	0.40	-0.00 (0.84)	1829	0.842
Lives with parents	0.52	0.50	0.02 (0.52)	1829	0.521
Any permanent work experience	0.13	0.34	-0.01 (0.73)	1829	0.730
Searched for work (last 6 months)	0.75	0.43	0.01 (0.83)	1829	0.832
Age	23.44	3.00	0.22 (0.23)	1829	0.230
Years since school	42.30	273.93	-10.95 (0.49)	1826	0.492
Search frequency (weeks of last 2 months)	0.57	0.31	0.00 (0.89)	1829	0.889
Work frequency (weeks of last 2 months)	0.34	0.38	-0.01 (0.61)	1829	0.611
Self employed	0.05	0.22	0.01 (0.60)	1829	0.601
Casual labourer	0.06	0.23	-0.02	1829	0.087

Satisfied with job	0.09	0.28	(0.09) -0.01	1829	0.659
Total savings	2279.23	6203.56	(0.66) 290.89	1829	0.346
Reservation wages	1327.22	1235.30	(0.35) 34.35	1808	0.632
Distance from city centre (km)	5.92	2.24	(0.63) -0.60	1829	0.229
Trips to the city centre (7d)	1.83	2.03	(0.23) 0.21	1826	0.185
Has formal job	0.06	0.23	(0.19) 0.00	1829	0.810
Uses CV in applications	0.28	0.45	(0.81) -0.00	1829	0.903
Expected no. job offers	1.46	2.09	(0.90) -0.21	1697	0.245
Aspired wage	5583.33	5830.85	(0.24) 191.89	1694	0.636
No. job contacts	6.74	9.63	(0.64) 0.89	1818	0.529
Present biased	0.12	0.33	(0.53) 0.00	1252	0.889
Future biased	0.08	0.27	(0.89) -0.02	1252	0.282
Life satisfaction	4.20	1.85	(0.28) -0.08	1828	0.633
			(0.63)		

Note: This table reports our baseline balance tests. For each baseline outcome of interest, we report the p-values for a test of the null hypothesis that we have balance between treatment and control groups. We cannot reject the null for any of the variables.

Appendix Table B.2: Summary of variables used in blocking/re-randomisation

	N	Mean	S.Dev.	1st Q.	Median	3rd Q.	Min.	Max.	p-value
Private limited company	493	0.51	0.50	0.00	1.00	1.00	0.0	0.0	1.0
NGO	493	0.13	0.34	0.00	0.00	0.00	0.0	0.0	1.0
Tours & Hospitality	493	0.19	0.39	0.00	0.00	0.00	0.0	0.0	1.0
Services & Finances	493	0.21	0.41	0.00	0.00	0.00	0.0	0.0	1.0
Education & Health	493	0.21	0.41	0.00	0.00	0.00	0.0	0.0	1.0
Manufacturing	493	0.26	0.44	0.00	0.00	1.00	0.0	0.0	1.0
Construction & Mining	493	0.14	0.35	0.00	0.00	0.00	0.0	0.0	1.0
Distance to centre	491	4.93	8.85	1.96	3.42	5.80	0.2	123.6	0.886
Total employees	493	288.11	972.98	37.00	87.00	225.00	4.0	18524.0	0.598
Workforce composition (job category)									
Professionals	493	0.29	0.23	0.10	0.21	0.45	0.0	0.0	0.9
Support sta	493	0.24	0.15	0.13	0.22	0.32	0.0	0.8	0.401
Production	493	0.26	0.29	0.00	0.17	0.50	0.0	1.0	0.863
Customer services	493	0.14	0.16	0.00	0.07	0.22	0.0	0.7	0.873
Workforce composition (education)									
Degree	493	0.23	0.24	0.04	0.13	0.37	0.0	1.0	0.901
Diploma	493	0.17	0.15	0.05	0.13	0.24	0.0	1.0	0.519
Turnover	493	0.21	0.88	0.05	0.10	0.19	0.0	14.3	0.150
Total annual new years	493	54.45	218.42	4.00	11.00	35.00	0.0	3901.0	0.268
Hiring rate	492	0.31	0.97	0.05	0.13	0.26	0.0	14.3	0.433
Use formal recruitment	493	0.65	0.48	0.00	1.00	1.00	0.0	1.0	0.703
Would come to a fair	493	0.79	0.41	1.00	1.00	1.00	0.0	1.0	0.711
Total sales (1000s)	339	554.75	3.84e+03	7.1750	23.017	121.8310	0.0	6.0e+04	0.492
Average salary (Birr)	493	2885.07	3010.35	1303.03	1990.18	3190.00	0.0	27683.2	0.812
Expected hiring rate	493	0.22	0.85	0.00	0.08	0.19	0.0	14.9	0.571

Notes: This table provides basic descriptive statistics on sample rms; in doing so, it also shows the variables used for blocking and re-randomisation. The 'p-value' column shows individual p-values for tests of covariate balance.

Appendix Table B.3: Worker employment amenities

Outcome	Estimated ITT	Control Mean	Observations
Received job by interview	0.0270 (.141) [1]	0.167	1702
Office work (7d)	0.00700 (.803) [1]	0.201	1702
Skills match with tasks	-0.0380 (.219) [1]	0.130	1702
Overqualified	0.0290 (.395) [1]	0.291	1702
Underqualified	-0.0130 (.468) [1]	0.0820	1702

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q -values are reported in square brackets.

Appendix Table B.4: Worker job search outcomes

Outcome	Estimated ITT	Control Mean	Observations
Applied to temporary jobs	0.242 (.347) [.533]	1.311	1693
Applied to permanent jobs	-0.0670 (.749) [.713]	2.279	1692
Interviews/Applications	0.0190 (.539) [.706]	0.354	972
Offers/Applications	-0.00300 (.937) [.881]	0.248	975
Interviews/Applications (Perm)	0.0850 (.039)** [.365]	0.327	742
Offers/Applications (Perm)	0.0790 (.114) [.365]	0.164	742
Interviews/Applications (Temp)	-0.0680 (.08)* [.365]	0.389	586
Offers/Applications (Temp)	-0.0630 (.207) [.401]	0.332	586
Uses CV for applications	-0.0530 (.074)* [.365]	0.401	1702
Uses certificates	0.0180 (.711) [.713]	0.479	1702

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q -values are reported in square brackets.

Appendix Table B.5: Median rate of expected number of new hires in the coming 12 months, as a percentage of current workforce

Industry	Worker Type					All workers
	Client services	Production	Support sta	White collar		
Construction, Mining, Farming	0.0%	14.3%	9.2%	15.4%		20.0%
Tours-Hospitality	16.7%	10.8%	10.2%	10.6%		14.8%
Finance, Services, Retail	10.5%	6.3%	10.1%	16.0%		16.1%
Education, Health, Aid	4.5%	5.7%	5.0%	14.3%		13.0%
Manufacturing	0.0%	8.0%	1.6%	3.4%		8.8%
All Industries	7.4%	9.3%	7.4%	11.1%		12.6%

Note: This table shows rms' stated expectations about new hires in the coming 12 months.

Appendix Table B.6: Correlates of worker attendance at the job fairs

	(1) Background	(2) Search Effort	(3) Employment	(4) All
Degree	0.0639 (0.198)			0.0330 (0.209)
Vocational	0.00802 (0.0395)			0.00559 (0.0398)
Post_secondary	0.000127 (0.191)			-0.0294 (0.201)
Female	-0.0109 (0.0307)			-0.0115 (0.0310)
Migrant	0.0154 (0.0362)			-0.00141 (0.0358)
Amhara	0.00957 (0.0376)			0.0148 (0.0338)
Oromo	-0.0181 (0.0506)			-0.0164 (0.0488)
Experience	-0.0590 (0.0547)			-0.0433 (0.0533)
Age	-0.00861 (0.00528)			-0.00924* (0.00518)
Certificate	0.0984*** (0.0304)			0.0654* (0.0357)
Distance (center)	0.00214 (0.00722)			0.00167 (0.00715)
Search_6months		0.0418 (0.0409)		0.0155 (0.0469)
Plan Self Empl		0.0399 (0.0898)		0.0297 (0.0891)
Search frequency		0.304*** (0.0497)		0.293*** (0.0505)
Wage Empl (6 months)			-0.0164 (0.0304)	-0.0446 (0.0289)
Work frequency			-0.0291 (0.0496)	-0.00877 (0.0524)
Employment at the time of the job fair				
Permanent Job			-0.161** (0.0646)	-0.160** (0.0692)
Any Job			-0.00143 (0.0338)	-0.00576 (0.0335)
Constant	0.748*** (0.253)	0.398*** (0.0376)	0.631*** (0.0270)	0.664** (0.263)
Observations	1,006	1,006	1,006	1,006
R-squared	0.018	0.045	0.007	0.063

Note: This table reports regression coefficients from a Linear Probability Model, in which we regress attendance at the fairs on a rich set of baseline characteristics; we provide robust standard errors in parentheses. We find no evidence suggesting that observably weaker candidates are more likely to attend the fairs: education, gender, and current employment do not significantly predict attendance. The only two robust predictors of attendance are instead associated with a positive motivation to work: attendance is higher among those job-seekers who search the most at baseline and who produce a formal certificate to employers.

Appendix Table B.7: Main industry classifications

Main Industry	Frequency	Percent
Tours-Hospitality	92	18.7
Finance, Services, Retail	102	20.7
Education, Health, Aid	104	21.1
Manufacturing	126	25.6
Construction, Mining, Farming	69	14.0
Total	493	100

Note: This table shows the initial partitioning of firms into five main industries prior to randomisation.

Appendix Table B.8: Blocking variables for the rm randomisation

variable	definition	source (question number)
plc	Firm is a private limited company	g3 = 3
total_n_all	Total number of pay-roll employees at the rm	l1_1_n
prop_p	Proportion of workers who are professionals	l1_5_n/l1_1_n
ed.deg	Number of workers at the rm with a degree	rowtotal (l1_19_1)=rowtotal(ed_total)
to_all	Rate of turnover in the last year	rowtotal (l2_1_)=total_n_all
formal_adv	Firms advertise when recruiting for jobs	l4_2_1=1 or l4_2_2=1
fairs	Firms expressed interest in attending a job fair	l4_31
hire_all	Rate of new hiring in the last year	rowtotal (l3_2_)=total_n_all

Note: This table defines the variables used for blocking for rm randomisation.

Appendix Table B.9: Correlates of firm attendance at the job fairs

	(1) Blocking	(2) Others	(3) Salaries	(4) All
Tours-Hospitality	-0.210*			-0.742**
	(0.117)			(0.351)
Finance, Services, Retail	-0.0150			-0.244
	(0.119)			(0.347)
Education, Health, Aid	-0.105			-0.674
	(0.130)			(0.652)
Manufacturing	-0.0556			-0.425
	(0.108)			(0.301)
Distance from city centre (km)	0.00270			0.0352
	(0.00385)			(0.0231)
Total employees (100s)	0.00171			-0.00377
	(0.00586)			(0.0203)
Respondent is owner	0.0306			0.0573
	(0.0869)			(0.251)
Turnover Rate	-0.0600			1.343
	(0.223)			(1.505)
Quit rate	-0.0268			0.453
	(0.252)			(1.799)
Workers with degrees	-0.427**			-0.772
	(0.197)			(0.912)
Workers with highschool	-0.0534			0.962**
	(0.174)			(0.456)
Proportion professionals	0.0114			1.611*
	(0.228)			(0.922)
Proportion female	0.144			0.460
	(0.175)			(0.397)
Total sales (log)		-0.0377		-0.0578
		(0.0340)		(0.0628)
Hiring Rate		0.248		-0.633
		(0.304)		(0.595)
Number permanent hires		0.0686		0.166
		(0.142)		(0.154)
Employee growth rate		-1.477		-2.275
		(1.347)		(1.765)
Growth rate (professionals)		0.120		0.704
		(0.437)		(0.500)
Growth rate (service)		0.0176		0.289*
		(0.137)		(0.157)
Growth rate (production)		0.917		1.122
		(0.689)		(0.947)
Growth rate (support)		0.0536		-0.309
		(0.366)		(0.414)
Starting salaries (professionals)			-0.0517	-0.106
			(0.192)	(0.260)
Starting salaries (services)			0.279	0.204
			(0.184)	(0.354)
Starting salaries (production)			0.163	0.254
			(0.187)	(0.303)
Starting salaries (support)			-0.142	-0.181
			(0.214)	(0.272)
5 year salary (professionals)			-0.116	0.0375
			(0.207)	(0.278)
5 year salary (services)			-0.0966	-0.328
			(0.224)	(0.321)
5 year salary (production)			-0.169	-0.228
			(0.195)	(0.266)
5 year salary (support)			0.0915	0.367
			(0.196)	(0.284)
Constant	0.834***	1.051**	1.302	0.835
	(0.128)	(0.411)	(0.987)	(1.465)
Observations	232	70	87	61
R-squared	0.075	0.075	0.102	0.576

Note: This table reports results from a series of Linear Probability Models; in each case, the outcome variable is a dummy for whether a firm attended the job fairs, conditional upon having been invited. Parentheses show heteroskedasticity-robust standard errors. The omitted industry dummy is for 'construction/mining'.

Appendix Table B.10: Determinants of attrition among job-seekers

Fairs	-0.025** (0.012)	Oromo	-0.007 (0.016)
Work frequency (weeks of 2 months)	0.007 (0.018)	Wage empl (6m)	0.017 (0.014)
Degree	-0.024 (0.017)	Married	-0.015 (0.017)
Worked (7d)	-0.015 (0.016)	Years since school	0.000 (0.0027)
Searched job (7d)	0.008 (0.014)	Lives with parents	0.008 (0.015)
Female	0.029** (0.013)	Ever had permanent job	0.002 (0.019)
Respondent age	0.000 (0.0027)	Searched job (6m)	-0.020 (0.017)
Born outside Addis	0.031** (0.015)	Amhara	0.000 (0.014)
		Constant	0.061 (0.060)
Average Attrition	6.7%	R-squared	0.012
Observations	1,827	F-test (treatment)	4.320
F-test (covariates)	1.130	p-value (treatment)	0.038
p-value (covariates)	0.320		

Note: This table reports regression results from a Linear Probability Model, in which the dependent variable is a dummy for whether a job-seeker attrited between baseline and endline; parentheses show heteroskedasticity-robust standard errors.

Appendix Table B.11: Firm recruitment in the last year

<i>Outcome</i>	Estimated ITT	Control Mean	Observation
<i>Panel A: Short term recruitment outcomes</i>			
Time taken to fill professional vacancies	-2.344 (1.986) [.658]	24.11	338
Time taken to fill non-professional vacancies	0.724 (1.751) [.909]	15.66	109
Number of interviews per position (professional)	0.312 (2.355) [.909]	8.818	361
Pay per recruitment (professional)	746.7 (1030.791) [.909]	2818	382
Pay per recruitment (non-professional)	-437.8 (320.543) [.658]	1259	406
Unfilled vacancies	0.601 (.247)** [.101]	0.859	305
<i>Panel B: Characteristics of workers recruited</i>			
Number of new hires for the year (professional)	-1.604 (2.688) [1]	11.73	472
Number of new hires for the year (non-professional)	-9.704 (7.283) [1]	44.64	472
Did firms mostly hire people with degrees (professional positions)?	-0.00800 (.041) [1]	0.574	473
Percentage of new hires hired in permanent positions (non-professional)	-0.00900 (.03) [1]	0.892	337
Percentage of new hires hired in permanent positions (professional)	-0.00800 (.031) [1]	0.876	308

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets, corrected for the tests conducted within each panel.

Appendix Table B.12: Impacts on firm hiring after job fairs

<i>Outcome</i>	Estimated ITT	Control Mean	Observations
Number of vacancies	0.169 (.266) [1]	1.115	422
New Hires	-0.671 (.866) [1]	3.907	422
Hiring shortfall	-0.0160 (.034) [1]	0.0290	193
Unfilled vacancies	0.380 (.785) [1]	2.143	422

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets.

Appendix Table B.13: Impacts on firm hire quality after job fairs

<i>Outcome</i>	Estimated ITT	Control Mean	Observations
Permanent workers hired	0.0200 (.049) [1]	0.336	422
Days taken to recruit for position (avg)	0.311 (1.386) [1]	11.75	190
Starting salary of new recruits (avg)	-673.9 (636.454) [1]	1031	160
Workers with degrees hired (%)	-0.0430 (.044) [1]	0.237	422

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets.

Appendix Table B.14: Firms' total workforce composition

<i>Outcome</i>	Estimated ITT	Control Mean	Observations
Total number of employees	-18.38 (16.581) [.847]	350.5	473
Proportion of professional workers on permanent contracts	0.0190 (.019) [.847]	0.908	462
Proportion of non-professional workers on permanent contracts	0.0280 (.02) [.67]	0.896	408
Average starting salary (professional)	-53.52 (235.925) [1]	4280	454
Average starting salary (non-professional)	102.9 (126.66) [.847]	1059	400
Proportion of professional workers with degree	-0.0570 (.027)** [.366]	0.645	461
Proportion of workers with post-secondary education (non-professionals)	0.0370 (.027) [.67]	0.355	407
Average worker is not under-qualified in any of the worker categories	0.00300 (.038) [1]	0.752	473

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets.

Appendix Table B.15: Impacts on firm turnover and employee growth

<i>Outcome</i>	Estimated ITT	Control Mean	Observations
Firing rate (professionals)	0.00400 (.004) [1]	0.00600	458
Firing rate (non-professionals)	0.00300 (.005) [1]	0.0130	319
Quit rate (professionals)	0.00800 (.02) [1]	0.143	458
Quit rate (non-professionals)	0.0250 (.037) [1]	0.134	320
Employee growth rate	0.0170 (.016) [1]	0.0140	472
Employee growth rate (professionals)	-0.0140 (.03) [1]	0.0310	467

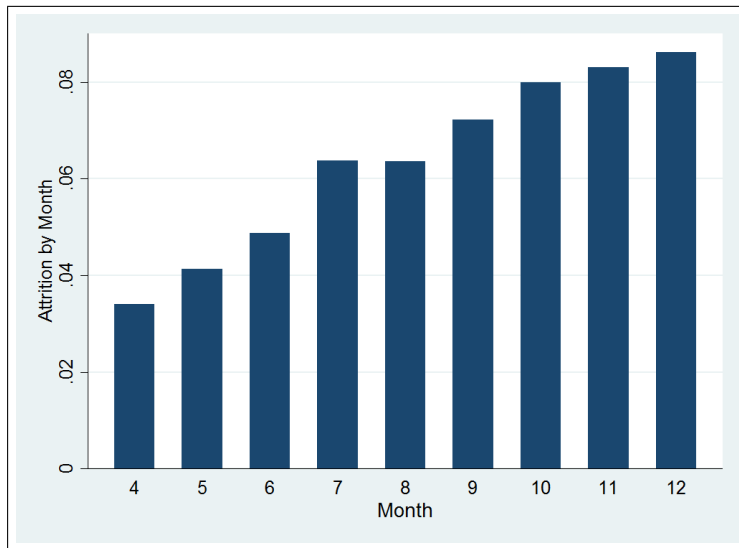
Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets.

Appendix Table B.16: Impacts on firm human resources policies and attitudes

<i>Outcome</i>	Estimated ITT	Control Mean	Observations
Firm reports HR problem	0.0820 (.037)** [.217]	0.752	473
Uses incentives in HR	0.0390 (.043) [.588]	0.595	473
Firm estimate of a fair wage	201.2 (312.897) [.592]	5463	452
Uses short term contractors	0.0480 (.045) [.588]	0.479	473
Uses performance rewards (professionals)	-0.0300 (.045) [.592]	0.545	473
Uses performance rewards (non-professionals)	-0.0740 (.045)* [.417]	0.562	473
Retrains poor performers	0.0390 (.04) [.588]	0.719	473

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets.

Appendix Figure B.1: Attrition rate from the Phone Survey by Month



Notes: This figure shows the trajectory of monthly attrition rates over the course of the phone survey. Attrition is defined as failure to complete one interview.

Appendix Table B.17: Impacts on firm growth and productivity

<i>Outcome</i>	Estimated ITT	Control Mean	Observations
Firm is for-profit	-0.0140 (.011) [1]	0.867	471
Sales Revenue (last year)	-17575 (23388.044) [1]	144370	331
Value Added	-15491 (11969.701) [1]	80851	327
Profit (inferred)	6026 (4791.574) [1]	12975	326
Self-reported profit	1853 (7175.053) [1]	29626	313
Capital stock	60034 (123774.721) [1]	185398	279
Investment (12 months)	-6452 (5920.8) [1]	20147	398
Sales per worker	-57.12 (76.278) [1]	604.5	330
Value added per worker	19.45 (28.102) [1]	220.3	326

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets.

Appendix Table B.18: **Firms' reasons for not hiring workers they met at the fairs**

Main self-reported reason	Percent
Insufficient work experience	34.38
Wrong expertise	7.03
Wrong educational qualifications	23.44
Poor performance at the interview	7.03
The candidates we wanted were hired by other firms	3.91
Poor references	2.34
Salary disagreement	2.42
Workers were not interested or did not apply	1.61
Workers arrived late	1.61
Firm did not have vacancies at the time of the fair	3.23
Other	10.48

Appendix Table B.19: Dyadic regressions: Firm requests to meet workers as function of worker characteristics

	(1)	(2)	(3)	(4)
	Firm requested to meet worker			
Worker has some permanent work experience	0.0173*** (0.00657)		0.0151** (0.00652)	0.0132** (0.00617)
Worker is recent graduate	0.00185 (0.00379)	0.00185 (0.00379)	-0.000304 (0.00450)	0.00153 (0.00479)
Worker has certificate with application	0.00190 (0.00273)	0.00190 (0.00273)	0.00136 (0.00285)	0.00182 (0.00286)
Worker has postsecondary education	0.00729*** (0.00257)	0.00729*** (0.00257)	0.00783*** (0.00270)	0.00808*** (0.00279)
Permanent work experience * fresh graduate	-0.00610 (0.00948)	-0.00610 (0.00948)	-0.00317 (0.01000)	-0.000833 (0.0107)
Permanent work experience * Highschool only	-0.0180*** (0.00663)	-0.000785 (0.00421)	-0.0162** (0.00679)	-0.0139** (0.00652)
Permanent work experience * postsecondary education		0.0173*** (0.00657)		
GS- algorithm suggested match			0.0256*** (0.00805)	0.0266*** (0.00833)
GS- matches we randomly suggested			-0.000233 (0.00792)	-0.00527 (0.00661)
Controls: Firms' vacancy characteristics	No	No	Yes	Yes
Controls: Firm baselien characteristics	No	No	No	Yes
Observations	19,110	19,110	18,185	17,491
R-squared	0.003	0.003	0.005	0.007

Notes: We regress on worker- firm dyadic data for all workers and firms who were invited to the same job fair, whether the firm requested to meet that worker in person, using a centralized meeting-request system facilitated at the job fairs. We include controls for worker characteristics, firm characteristics, and vacancy characteristics (vacancies held by the firm in question at the time of the job fairs).

Appendix Table B.20: Heterogeneous effects: rm recruitment by baseline proportion of youth in workforce

	(1)	(2)	(3)	(4)	(5)
	Firm performed formal interviews (professionals)	Firm performed formal interviews (non-professionals)	Did any advertising for new hires	Did advertising for professional positions	Did advertising on the job boards
Proportion young high	0.0121 (0.0356)	-0.0501 (0.0568)	0.0545 (0.0461)	0.0492 (0.0558)	0.0846 (0.0607)
Proportion young low	0.0689 (0.0534)	0.0123 (0.0546)	0.0554 (0.0443)	0.177*** (0.0535)	0.0959 (0.0583)
Observations	473	473	473	473	473
R-squared	0.009	0.013	0.017	0.027	0.010
ControlMean 1	0.715	0.659	0.805	0.634	0.325
ControlMean 2	0.647	0.555	0.773	0.555	0.336
Test 1=0 (p)	0.462	0.430	0.989	0.0997	0.893

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets, corrected for the tests conducted within each panel.

Appendix Table B.21: Heterogeneous effects: firm recruitment by baseline use of formal recruitment methods

	(1)	(2)	(3)	(4)	(5)
	Firm performed formal interviews (professionals)	Firm performed formal interviews (non-professionals)	Did any advertising for new hires	Did advertising for professional positions	Did advertising on the job boards
Used formal methods	0.0666 (0.0473)	0.00436 (0.0486)	0.0781** (0.0393)	0.155*** (0.0475)	0.0827 (0.0517)
Did not use formal methods	-0.0159 (0.0661)	-0.0737 (0.0679)	-0.00102 (0.0550)	0.0309 (0.0664)	0.103 (0.0721)
Observations	473	473	473	473	473
R-squared	0.004	0.003	0.008	0.023	0.010
ControlMean 1	0.776	0.671	0.888	0.714	0.410
ControlMean 2	0.494	0.481	0.593	0.358	0.173
Test 1=0 (p)	0.311	0.350	0.242	0.128	0.816

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets, corrected for the tests conducted within each panel.

Appendix Table B.22: Heterogeneous effects: firm workforce composition at endline by proportion of workforce in professional occupations

	(1) Firm performed formal interviews (professionals)	(2) Firm performed formal interviews (non-professionals)	(3) Did any advertising for new hires	(4) Did advertising for professional positions	(5) Did advertising on the job boards
Proportion professional high	-0.0223 (0.0542)	-0.0331 (0.0557)	0.0573 (0.0452)	0.105* (0.0546)	0.0530 (0.0592)
Proportion professional low	0.100* (0.0544)	-0.0110 (0.0560)	0.0452 (0.0454)	0.121** (0.0549)	0.127** (0.0595)
Observations	473	473	473	473	473
R-squared	0.007	0.001	0.005	0.018	0.011
ControlMean 1	0.788	0.559	0.822	0.678	0.356
ControlMean 2	0.581	0.653	0.758	0.516	0.306
Test 1=0 (p)	0.111	0.780	0.849	0.829	0.380

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets, corrected for the tests conducted within each panel.

Appendix Table B.23: Heterogeneous effects: firm hiring after the fair by proportion of workforce in professional occupations

	(1) Num. new hires (professional)	(2) Num. new hires (non-prof)	(3) Most hires with degrees (professionals)	(4) Percentage new hires perm. positions (non-prof)	(5) Percentage new hires perm positions (prof)
Proportion professional high	-4.761 (3.825)	-17.89* (10.35)	-0.0815 (0.0580)	0.00427 (0.0402)	-0.0210 (0.0478)
Proportion professional low	0.422 (3.836)	1.430 (10.38)	0.0624 (0.0583)	-0.0214 (0.0447)	0.00920 (0.0409)
Observations	472	472	473	337	308
R-squared	0.003	0.006	0.007	0.001	0.001
ControlMean 1	20.51	62.43	0.669	0.840	0.892
ControlMean 2	3.444	27.85	0.484	0.958	0.864
Test 1=0 (p)	0.339	0.188	0.0800	0.669	0.633

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets, corrected for the tests conducted within each panel.

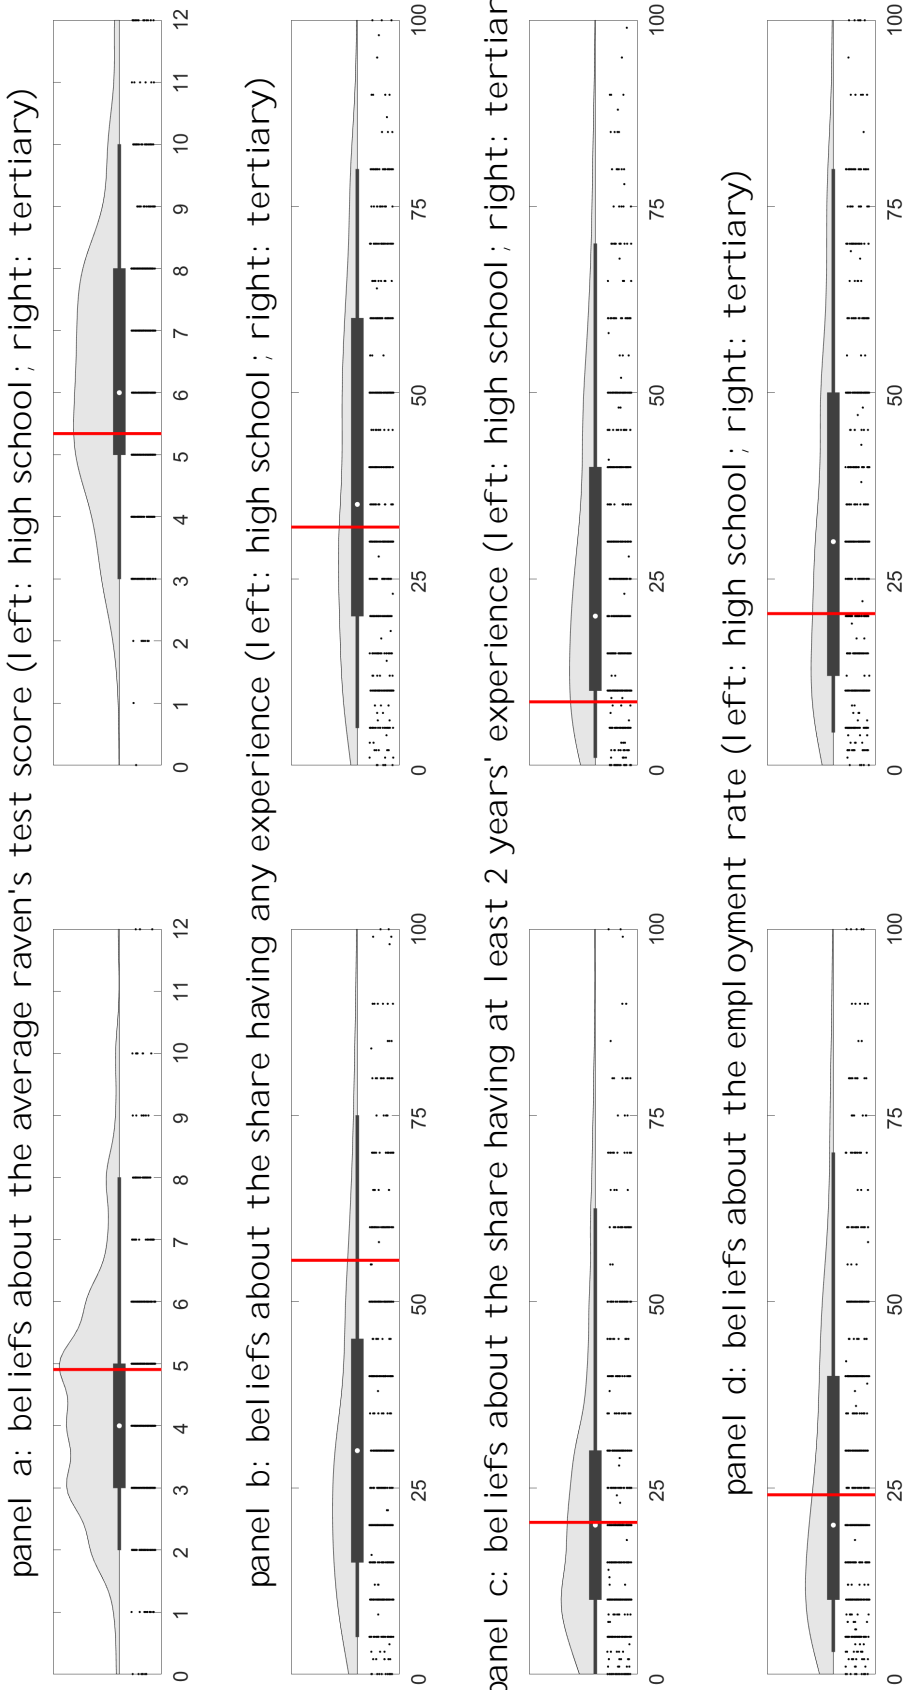
Appendix Table B.24: Heterogeneous effects: rm recruitment by proportion of workforce in professional occupations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Total number of employees	Prop. professional workers on perm. contract	Prop. non-prof workers on perm. contract	Av. starting salary (professionals)	Av. starting salary (non-prof's)	Prop. prof workers with degree	Prop. non-prof workers with degree	Average worker not under-qualified
Prop. professional high	-46.96** (23.54)	0.0352 (0.0272)	0.0438 (0.0302)	181.1 (338.0)	180.7 (188.4)	-0.0698* (0.0375)	0.0824** (0.0397)	-0.0262 (0.0536)
Prop. professional low	14.86 (23.66)	0.00629 (0.0274)	0.0180 (0.0275)	-187.9 (340.1)	63.95 (170.4)	-0.0427 (0.0380)	0.00720 (0.0363)	0.0331 (0.0538)
Observations	473	462	408	454	400	461	407	473
R-squared	0.009	0.004	0.006	0.001	0.003	0.010	0.011	0.001
ControlMean 1	501.5	0.862	0.881	4818	653.6	0.728	0.502	0.814
ControlMean 2	206.8	0.951	0.908	3760	1398	0.564	0.231	0.694
Test 1=0 (p)	0.0641	0.454	0.528	0.442	0.646	0.612	0.163	0.435

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

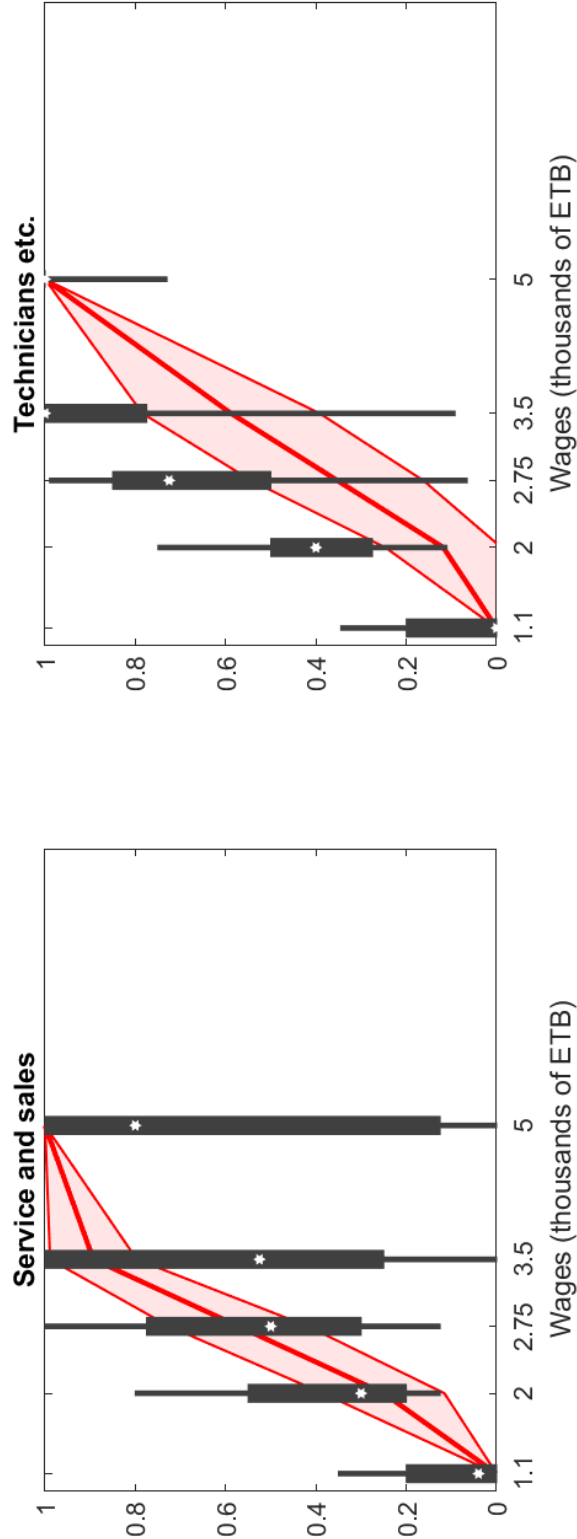
Notes: Each row reports a separate regression. For each regression, we report the estimated ITT from participating in the job fair, the mean in the control group, and the number of observations. Standard errors are reported in parentheses; False Discovery Rate q-values are reported in square brackets, corrected for the tests conducted within each panel.

Appendix Figure B.2: Distribution of employer beliefs by education



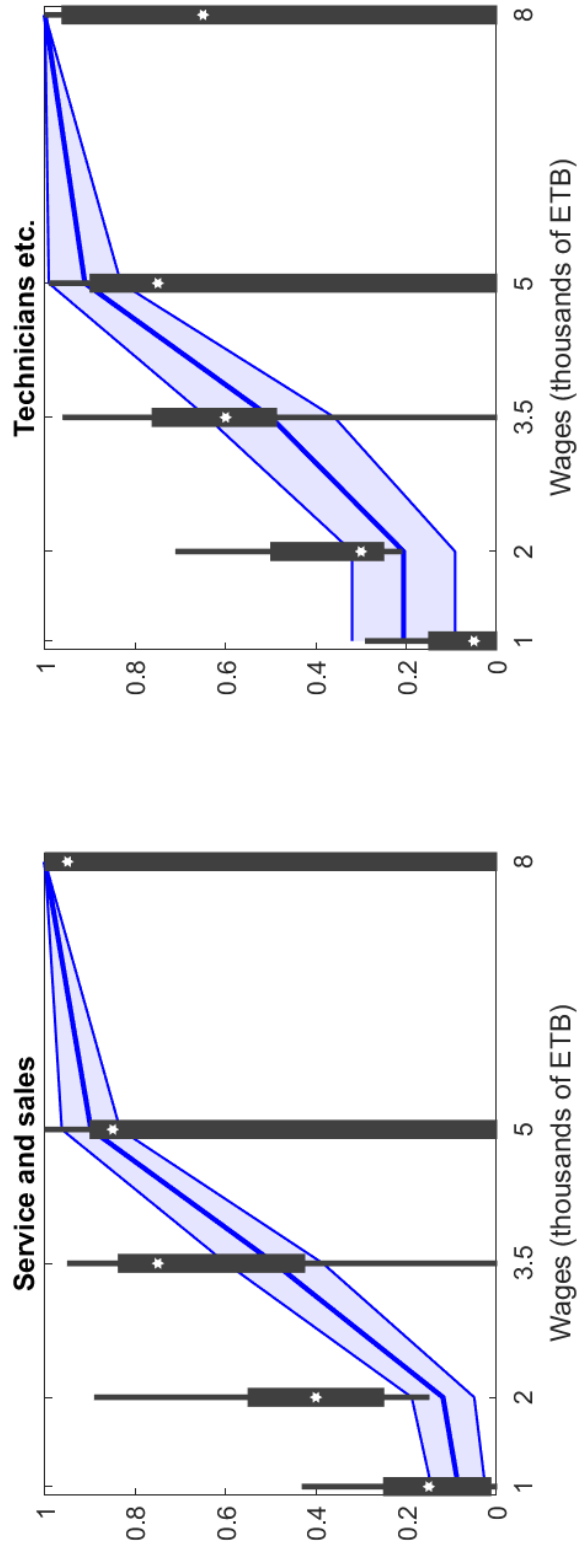
Notes: Each panel shows the distribution of employer beliefs, for some specific job-seeker characteristic. Each panel shows a kernel density plot (with bandwidth chosen using Silverman's plug-in), then a box-plot (showing 5th, 25th, 50th, 75th and 95th percentiles), followed by a direct visualisation of data (with random vertical jitter). For each panel, we superimpose a vertical line at the true value of the statistic in our sample.

Appendix Figure B.3: Firms' beliefs about the distribution of job-seekers' reservation wages: Secondary-educated job-seekers



Notes: These figures show the distribution of firms' beliefs about job-seekers' reservation wages. In each graph, we show this at five wage points: in each case, we show the distribution of firms' beliefs (using thin bars to show the 10th and 90th percentiles, thick bars to show the 25th and 75th percentiles, and a star symbol to show the median). The coloured lines show the true proportion of our job-seekers with a given reservation wage (where the shaded area represents the 90% confidence interval for the proportion).

Appendix Figure B.4: Job-seekers' beliefs about the distribution of firm wages: Secondary-educated job-seekers



Notes: These figures show the distribution of job-seeker beliefs about firms' wages. In each graph, we show this at five wage points: in each case, we show the distribution of job-seeker beliefs (using thin bars to show the 10th and 90th percentiles, thick bars to show the 25th and 75th percentiles, and a star symbol to show the median). The coloured lines show the true proportion of our firms with a given wage (where the shaded area represents the 90% confidence interval for the proportion).

