

**Manuscript version: Author's Accepted Manuscript**

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

**Persistent WRAP URL:**

<http://wrap.warwick.ac.uk/148081>

**How to cite:**

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

**Copyright and reuse:**

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

**Publisher's statement:**

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: [wrap@warwick.ac.uk](mailto:wrap@warwick.ac.uk).

# The Selection of Talent

## Experimental and Structural Evidence from Ethiopia\*

Girum Abebe,<sup>†</sup> A. Stefano Caria,<sup>‡</sup> Esteban Ortiz-Ospina<sup>§</sup>

November 25, 2020

### Abstract

We study how search frictions in the labour market affect firms' ability to recruit talented workers. In a field experiment in Ethiopia, we show that an employer can attract more talented applicants by offering a small monetary incentive for making a job application. Estimates from a structural model suggest that the intervention is effective because the cost of making a job application is large, and positively correlated with jobseeker ability. We provide evidence that this positive correlation is driven by dynamic selection. In a second experiment, we show that local recruiters underestimate the positive impacts of application incentives.

---

\*We are grateful to Abi Adams, Arun Advani, Nava Ashraf, Oriana Bandiera, Vittorio Bassi, Gaurav Chiplunkar, Stefano DellaVigna, David Eil, Marcel Fafchamps, Erica Field, James Fenske, Simon Franklin, Douglas Gollin, Clément Imbert, Supreet Kaur, Philipp Kircher, Rocco Macchiavello, Jeremy Magruder, David McKenzie, Guy Michaels, Paul Niehaus, Barbara Petrongolo, Simon Quinn, Imran Rasul, Chris Roth, Yared Seid, Alemayehu Seyoum Taffesse, Chris Woodruff, Liam Wren-Lewis, Yanos Zylberberg and seminar audiences at NBER Summer Institute, BREAD/CEPR/STICERD/TCD (LSE), Warwick, Toulouse, Cornell Dyson, Sciences Po, PSE, McGill, IFS, NEUDC (Tufts), Mannheim, Oxford, Bristol, Essex, Royal Holloway, Cologne, Ohlstadt for helpful comments. We thank Ian Jewitt for invaluable discussion of the theoretical model. Koen Maskaant and Alemayehu Woldu provided outstanding research assistance. The project would not have been possible without the support of Rose Page, Simon Quinn, CSAE and EDRI. We acknowledge funding from the IGC (Project No. 1-VCC-VETH-VXXXX-32304) and the John Fell Fund. The trial was registered on the AEA RCT registry (AEARCTR-0001900).

<sup>†</sup>World Bank. Email: [gtefera@worldbank.org](mailto:gtefera@worldbank.org).

<sup>‡</sup>University of Warwick. [stefano.caria@warwick.ac.uk](mailto:stefano.caria@warwick.ac.uk), Web: [www.stefanocaria.com](http://www.stefanocaria.com).

<sup>§</sup>Our World in Data and University of Oxford. Email: [esteban@ourworldindata.org](mailto:esteban@ourworldindata.org).

Selection problems occur everywhere in society. Firms want to hire talented workers; universities need to attract competent students; welfare programs have to select poor recipients. A key insight that often guides the design of selection policies is that a costly application process can improve selection by discouraging the participation of unwanted candidates. This logic underpins ordeal mechanisms in welfare programs (Nichols and Zeckhauser, 1982; Alatas et al., 2016) and is often applied in recruitment and marketing (Ashraf et al., 2010; Bandiera et al., 2011; Alonso, 2018). In this paper, we show experimentally that the opposite is true in an important economic context — *decreasing* application costs significantly improves selection. Our results are driven by the fact that high-ability candidates face on average higher application costs. We quantify this positive correlation between applicant ability and application costs, provide evidence on the mechanism that generates it, and show how it can be leveraged for policy.

Our evidence comes from studying an employer who wants to attract talented workers for a clerical position in Addis Ababa, Ethiopia. There are two important features in this context. First, job search and application costs are high on average; formal jobs require an application in person, which is time consuming and often requires the use of public transport. Second, these costs are heterogenous, as individuals have different access to liquidity and live at a varying distance from employers.

Using unique high-frequency data on job search and employment, we document that this heterogeneity generates a dynamic selection effect. High-ability individuals who face relatively low application costs find work faster and stop searching for work earlier than individuals who have similar ability, but face higher application costs. Thus, over time, a positive correlation between ability and application costs emerges among individuals searching for work.<sup>1</sup> This correlation overturns the standard intuition on the screening role of application costs and suggests that lowering costs may actually be a beneficial policy for an employer in this market. We test this prediction empirically with our experiment.

In the experiment, the employer reduces application costs by offering a small monetary payment to all job applicants. This monetary incentive is worth 4.5 USD and is calibrated to reimburse applicants for both transport costs — an in-person application is required for this position — and the opportunity cost of time. In a second treatment, the employer doubles the wage offer but does not provide any financial incentive for applications. The expected value of the wage increase is 105 USD. The employer ran-

---

<sup>1</sup>At the end of the paper, we present evidence from a number of studies suggesting that a similar correlation may be found in other contexts as well, in both developed and developing countries.

domises the offer of these two treatments over the sample of individuals who call to inquire about the position.

Our key finding is that the application incentive *improves* the quality of the applicant pool. We measure applicant ability through standardised tests of cognitive and non-cognitive ability, job-specific experience, and GPA (a measure of academic achievement widely used as a proxy of ability in this context).<sup>2</sup> Applicants from the application incentive group have higher cognitive ability and GPA compared to control applicants, and similar levels of non-cognitive ability and experience. The ability gains occur across the whole distribution of ability and are particularly large at the top. GPA and cognitive ability among the highest-scoring applicants invited for an interview increase by .5 and .3 standard deviations, respectively. The number of top applicants (defined as those with cognitive ability above the 90th percentile of the control group distribution) doubles. We also find that raising the wage increases average applicant ability. The magnitude of this effect is similar to the impact of the application incentive, although the two interventions operate through different mechanisms, as we discuss in detail below.

These findings surprise local employers. To show this, we sample 196 recruiters in the same market and ask them to predict the effect of the application incentive, after being informed of the characteristics of the applicant pool in the control and high wage conditions (DellaVigna and Pope, 2016). The majority of employers underestimate the effect of application incentives and the average employer incorrectly expects this intervention to decrease applicant quality. However, when asked to invite selected jobseekers to make an application at their firm on the basis of anonymised CVs, employers strongly prefer the applicants from the incentive treatment over the applicants from the other two experimental groups. This gives us further evidence that application incentives attract applicants that are preferred by employers.

The treatment effect on ability is driven by women, and by those jobseekers who are currently unemployed and less-experienced — all groups that do not usually perform well in the labour market. This is not the case for the high wage treatment. This suggests that the application incentive does not increase ability at the cost of attracting individuals who have better outside options and, hence, a lower likelihood of accepting and keeping the job. On the contrary, this intervention mostly taps the pool of low-income jobseekers who stand to benefit the most from the job. To explore this point further, we

---

<sup>2</sup>We use the Raven and Stroop tests for cognitive ability (Schmidt and Hunter, 1998). For non-cognitive skills we administer the Big-5 personality test and the Grit scale (John and Srivastava, 1999; Duckworth et al., 2007). To identify job-specific experience, we follow Autor and Handel (2013) and collect measures of experience in several relevant tasks.

generate an individual measure of the net present value of the experiment's job using a simple calibration framework and a Post-LASSO forecast of each individual's market wage (Belloni et al., 2014). We find that the increase in ability is significantly larger for the group of jobseekers that values the job the most.

We rule out several potential explanations for our findings that are not related to job application costs. First, the interventions could induce individuals to exert greater test effort. To study this, we administer a task that requires effort, but very little ability. We do not find significant differences in performance on this task, which provides evidence against differential test effort. The positive impacts on GPA, a measure of ability determined before the interventions, further support this conclusion. Second, the application incentive could help individuals overcome self-control problems. However, we find that incentive group applicants are as likely to be present biased as control group applicants, which is inconsistent with this explanation. Third, subjects could misinterpret the offer of the incentive as a signal about their quality or about the tightness of the labour market. Contrary to this hypothesis, we show that the incentive does not affect subjects' expectations about how long it will take them to find a new job, or the wage that this new job would pay. We also find that the incentive is associated with only minor changes in beliefs about the attributes of the job. We estimate that these changes can account for only 5 percent of the total effect of the application incentive.

To shed light on the structural features of the labour market that drive our reduced-form results we propose a simple model of application decisions. The model captures two key frictions in job search: application costs and uncertainty about the probability of being offered the job. There is only one type of vacancy, but we allow for worker heterogeneity — jobseekers differ in terms of their ability, the magnitude of the application costs they face, and the benefit that they derive from being offered the job. Using this model, we show formally that the incentive can attract better applicants in markets where higher-ability jobseekers face larger application costs.

We identify and structurally estimate the key parameters of the model using the exogenous variation generated by the experiment (DellaVigna, 2018). The simulated moments closely track the empirical moments used for estimation. For example, we fit all application rates with less than one percentage point of error. Further, the model can match a key non-targeted moment — jobseekers' assessment of the probability of receiving a job offer — and replicates non-targeted patterns of the data such as the fact that control applicants are negatively selected.

We estimate that application costs are large and, consistently with our initial descriptive evidence, positively correlated with jobseeker ability. For the group of indi-

viduals who value the job the most, the correlation between application costs and job-seeker ability is .57. The magnitude of application costs is also substantial. At the mean, application costs amount to 13.5 percent of the monthly wage for the same group of jobseekers. These central findings are robust to the use of different assumptions about the information available to jobseekers, alternative sets of empirical moments, and different restrictions on parameter heterogeneity.<sup>3</sup> Using a local estimate of the value of cognitive ability, we calculate that for the average firm in this market the internal rate of return (IRR) of the application incentive is 9.8 percent. This is much higher than the IRR of the high wage treatment, which is a costlier intervention. However, this IRR estimate is somewhat imprecise: the 90 percent bootstrapped confidence interval has a lower bound of -18.7 percent and an upper bound of 38.3 percent. Through counterfactual policy analysis, we show that the IRR increases substantially and becomes significantly different from zero when the incentive is either (i) targeted to the demographics that drive the treatment effects (e.g. women) or (ii) offered conditional on a good performance on the selection test.

Our results make several contributions to the literature. First, we highlight that upfront costs can worsen selection in an important economic context. To our knowledge, ours is the first worker selection experiment that manipulates application costs. Some recent experiments in developing and developed countries have manipulated the wage, or workers' expectations about the wage (Dal Bó et al., 2013; Deserranno, 2019; Belot et al., 2017; Ashraf et al., 2018). These studies typically find that higher wages attract better applicants. For example, Ashraf et al. (2018) show that offering career incentives enables the Zambian government to recruit more talented nurses.<sup>4</sup> A number of papers set in the US have also studied how various contract features affect applications decisions (Flory et al., 2014; Mas and Pallais, 2016). Finally, a recent paper by Hardy and McCasland (2017) studies the experimental placement of apprentices in small firms in Ghana and finds evidence consistent with hiring frictions. None of these studies directly varies applications costs. Our findings highlight that, when jobseekers find it costly to participate in the labour market, firms may hire better workers if they reduce application costs.

Second, we contribute to a recent, growing literature that studies frictions in the

---

<sup>3</sup>A potential concern is that unobserved variation in the value of the job may inflate the estimates of application costs and of the cost-ability correlation. While we cannot fully rule out this possibility, we show that our key results are robust to the use of two different estimation strategies that plausibly reduce this unobserved variation.

<sup>4</sup>An exception is Deserranno (2019), who finds that higher expected salaries select less motivated candidates for a non-profit organisation in Uganda.

allocation of talent. Previous studies have focused on the role of discrimination (Hsieh et al., 2013), migration costs (Bryan and Morten, 2015; Imbert and Papp, 2016; Lagakos et al., 2017), housing market failures (Hsieh and Moretti, 2015), and corruption (Weaver, 2016). We provide original empirical evidence on the importance of search frictions — in particular, frictions in the job-application process. These frictions have been the focus of several theoretical papers, but direct evidence on their magnitude has been limited to date (Banerjee and Newman, 1993; Marimon and Zilibotti, 1999; Rogerson et al., 2005; Paserman, 2008; Galenianos et al., 2011). A unique feature of our study — and a key contribution to this literature — is that search frictions are identified using exogenous experimental variation.

Our findings are consistent with those of an emerging literature that studies spatial frictions in urban labor markets. This literature shows that across a range of different contexts — including the US, Ethiopia and South Africa — transport subsidies increase job search intensity and impact labor market outcomes (Phillips, 2014; Franklin, 2017; Abebe et al., 2020; Banerjee and Sequeira, 2020). These papers, however, do not observe the counterfactual workers that would have been attracted by firms in the absence of the program, and thus do not capture changes in the selection of talent.<sup>5</sup> Our study, on the other hand, documents that decreasing application costs enables employers to attract higher-ability applicants, by incentivising a pool of female, inexperienced, unemployed jobseekers that are unlikely to quickly secure good positions otherwise. This suggests that job search assistance policies may have positive impacts on the allocation of talent and motivates the implementation of new market-level evaluations designed to investigate these effects (Crépon et al., 2013).

Lastly, our results highlight that managers do not have accurate beliefs about the returns to different recruitment practices and may thus fail to optimise firms' recruitment policies. Providing information to managers may thus be a cost-effective intervention in this context. These results contribute to the nascent literature on behavioural firms (DellaVigna and Pope, 2016; DellaVigna and Gentzkow, 2017; Kremer et al., 2018).

## I Context

In this section we describe the labour market in Addis Ababa from the point of view of both firms and workers. We rely both on descriptive statistics from our data on jobseek-

---

<sup>5</sup>Crépon and Van den Berg (2016) and McKenzie (2017a) offer recent reviews of the job search assistance literature. Additionally, there is a relevant literature at the intersection of urban and labour economics that studies spatial frictions that result from urban segregation, mostly in developed countries, which is summarised in Gobillon et al. (2007) and Zenou (2009).

ers and firms, as well as on data from the longitudinal, high-frequency labour market survey collected by [Abebe et al. \(2020\)](#). We show that returns to cognitive ability are likely to be substantial in this context; that job search and application costs are an important barrier for many jobseekers; and that high-cost jobseekers may be better selected than low-cost jobseekers.

## **I.A The challenge of finding high-ability workers**

Finding a worker with the right ability and skills can be challenging. To collect data on employers in Addis Ababa and their beliefs, we sample 196 firms that advertised a vacancy for a clerical job during a period of six weeks in 2017 and ask managers about the HR problems they face and the HR practices they have adopted.<sup>6</sup> The most frequently mentioned HR challenge is finding workers with the right skills. As shown in [Figure 1](#), about 35 percent of managers consider this to be the most pressing HR problem for their firm. Retention, absenteeism, motivation and conduct are all mentioned less frequently than hiring. In terms of HR strategies, about 60 percent of managers report that offering higher wages is the most effective way to improve the quality of recruits. Application incentives are mentioned rarely and, in practice, they are not frequently used by firms in the city.

< **Figure 1 here.** >

The firms in our sample hire workers on a frequent basis. In the two months preceding the interview, the average monthly hiring rate among these firms was 2.5 percent and the average separation rate was 1.6 percent. Hiring thus occurs both to expand the workforce and to replace workers who leave the firm. These labour flows are somewhat smaller than those experienced by firms in the US (e.g. in the US, in June 2017, the average monthly hiring rate was 3.7 percent and the average monthly separation rate was 3.6 percent ([Bureau of Labor Statistics, 2017](#))).

We also find that hiring is costly for firms, in terms of both time and money. Among firms in our sample, average recruitment costs amount to about 104 USD and 18 hours

---

<sup>6</sup>The firms are selected in the following way. First, we screen all vacancies advertised on the main job-vacancy boards or in a popular newspaper insert. To identify clerical jobs, we categorise each vacancy according to the 2010 Standard Classification of Occupations of the US Bureau of Labor Statistics. For the full list of occupations included in the survey see [Table A1](#). In each firm, we request to interview the head of the selection committee – typically the head of the HR department or the firm’s CEO. We use this sample of managers to run the second experiment reported in the paper. During the interview, each manager first completes the CV-ranking and forecast tasks, which we describe in detail in [Section VI](#), and then answers the survey questions about his or her firm.

of staff time (worth about 40 USD when valued at the mean wage of an HR manager in the same firms). Total costs correspond approximately to one month of salary for one of the high-wage jobs in the experiment. These costs do not vary substantially with the number of applicants (many of the costs, such as those related to advertising and developing tests and interviews, are fixed). Managers estimate that considering one more application entails no further monetary costs and would not require more than one hour of staff time.

Firms usually screen workers by assessing CVs and by administering written tests and interviews. Educational qualifications, GPA and previous work experience are the most important variables that managers consider when they assess candidates' CVs. Firms often require applicants to deposit their CV and the other application materials in person. Written tests and interviews are also used frequently. Both interviews and written tests are used to assess general cognitive ability, specific technical knowledge, and personality traits.

## **I.B The cost of finding a job and the returns to ability**

From the perspective of jobseekers, finding a job is also challenging in this labour market. First, jobseekers spend substantial amounts of time and money to identify vacancies and apply for them. Using self-reported expenditure data, [Abebe et al. \(2020\)](#) estimate that the monetary cost of searching and applying for jobs amounts to one quarter of weekly expenditure for individuals who are actively looking for employment. To pay for these costs, jobseekers need to frequently take up informal, short-term jobs, which are easier to secure. These challenges are described in detail in [Abebe et al. \(2020\)](#). Here we report one additional piece of descriptive evidence: jobseekers apply for a very small fraction of the available vacancies. In our sample, for example, the average unemployed person in the control group completes approximately two job applications in 30 days. On the other hand, when we screened job boards and newspapers over a similar time period, we were able to find at least 30 relevant vacancies per week. This low number of applications is consistent with the existence of financial constraints that limit job search intensity.

Second, ability matters: highly talented individuals — particularly those with strong cognitive ability — earn more than low-ability individuals. In [Table A2](#), using data from [Abebe et al. \(2020\)](#), we show that a one-standard-deviation increase in a jobseeker's Raven test score is significantly associated with a 9 percent rise in wages, while similar improvements in conscientiousness or neuroticism are not significantly correlated with wages. Better performance on the Raven test is also significantly related to em-

ployment: a one-standard-deviation increase in the Raven score is associated with a 4.5 percentage point (8 percent) gain in the probability of employment. Including standard controls for age and labour market experience does not affect these results (Table A3).<sup>7</sup>

### I.C Are jobseeker ability and application costs correlated?

Lastly, we find evidence that ability and application costs are correlated among jobseekers.<sup>8</sup> Here, we provide descriptive evidence on this correlation, and on the dynamic selection mechanism that may drive it. In section V, we present experimental evidence consistent with these descriptive findings. Our hypothesis is that, in a given cohort, high-ability individuals who face low application costs find work faster and leave job search earlier compared to individuals who have similar ability but face higher application costs.<sup>9</sup> Thus, low-cost, high-ability individuals become progressively under-represented in the pool of jobseekers compared to high-cost, high-ability types. This generates a positive correlation between ability and application costs among jobseekers, which strengthens over time.

We provide evidence on this correlation using the longitudinal data collected by Abebe et al. (2020). This dataset is ideal for this exercise as it provides fortnightly data on the job-search decisions and employment outcomes of a sample of young adults in Addis Ababa for the period of one year. The sample was restricted to individuals who, at the start of data collection, did not have a formal, open-ended work contract. By the end of the year, about half of the people in the sample found employment. The data thus covers a period of active job search and job finding. The data also includes a Raven test administered close to the beginning of the panel and two variables which can proxy for search costs: a measure of financial resources (savings at baseline) and a proxy for transport costs (distance from the city centre).

We find three pieces of evidence that support our selection hypothesis. First, high-cost, high-ability jobseekers are less likely to find a job than low-cost jobseekers with

---

<sup>7</sup>For a systematic discussion on the returns to talent and, more broadly, human capital in developing countries see Porzio (2017) and Caselli et al. (2014).

<sup>8</sup>We use the term ‘jobseeker’ to indicate those individuals who are actively looking for employment at a given point in time. Thus, in our experimental sample, all individuals are jobseekers, as they have called to inquire about the experiment’s position — an active job-search step. In the sample of Abebe et al. (2020), on the other hand, we observe both individuals who look for employment and individuals who do not.

<sup>9</sup>The reverse may happen among low-ability individuals, who have a low chance of finding formal employment. Those who are low ability and high cost may decide to give up on the search for formal work altogether, and restrict themselves to the informal labour market instead. On the other hand, low-ability types who face low costs may invest in formal job search for a longer period of time.

equal ability. This is in contrast to low-ability individuals, for whom job-finding rates are similar irrespective of application costs. We show this point through the regression analysis reported in Table A4 in the Appendix, where we study month-to-month transitions from job search to employment. We find that, compared to low-cost types, high-cost individuals experience a lower increase in the probability of finding work for the same increase in ability — a significant effect when costs are proxied by savings and an insignificant effect of a roughly similar magnitude when costs are proxied by distance. Among jobseekers with ability one standard deviation above the mean, high-cost types have a probability of finding a job that is 8 percentage points lower than low-cost types. On the other hand, among jobseekers with ability one standard deviation below the mean, high and low-cost types have very similar job-finding rates (about one percentage point apart).<sup>10</sup>

Second, the correlation between jobseeker ability and application costs strengthens over the course of the year. In Figure 2 and Figure 3, we plot the fortnightly value of the cost-ability correlation in the selected sample of jobseekers. When costs are proxied by savings, we estimate that the correlation between jobseeker ability and application costs grows by a significant 0.011 of a standard deviation every fortnight. When costs are proxied by distance from the city centre, the correlation grows by a significant 0.018 of a standard deviation every fortnight. In other words, there is a growing gap in average ability between low and high-cost jobseekers (which we also illustrate showing the separate trends of high and low-cost types in Figure A1 in the Appendix).

< **Figures 2 and 3.** >

Third, by the end of the year, the correlation between jobseeker ability and application costs is large and positive. As we show in Figure 2, this correlation is large (and just marginally insignificant) for our first proxy of applications costs — low savings. Among the individuals who search for work in the last fortnight of the panel, those with below-median savings have a Raven test score that is .3 standard deviations higher than those with savings above the median. In the last fortnight of the panel, the correlation with our second proxy — distance from the city centre — is also positive, but it is smaller

---

<sup>10</sup>These figures refer to the estimates reported in column 1 of Table A4, where we proxy application costs by savings. Individuals with high savings may also have higher reservation wages. In the data, a one standard deviation increase in savings is associated with self-reported reservation wages that are about 0.15 standard deviations higher. Higher reservation wages should reduce the job finding rates of high saving individuals, partly cancelling the effect of their lower search costs. Thus, the results on job finding that we report here likely underestimate the differential selection that is due to search costs alone.

and less precisely estimated than the correlation with savings (Figure 3). Jobseekers who live above the median distance from the city centre have a Raven test score that is an insignificant .08 standard deviations higher than jobseekers who live below the median distance.

Overall, this evidence supports the hypothesis that selection dynamics generate a positive correlation between jobseeker ability and application costs. This observation motivates a model that explicitly considers how such a correlation can affect the hiring outcomes of a firm, which we present in the next section.

## II A simple model of job application decisions

We propose a simple model of application decisions that captures two key frictions in job search: application costs and uncertainty about the probability of being offered the job. The model describes the effects of application incentives on application rates and on the quality of the applicant pool.

**Jobseeker Characteristics.** Consider a set of individuals deciding whether to apply for the experiment's job. For tractability, let us focus on the large-number case and assume that these jobseekers form a continuum of unit measure.

Jobseekers differ in terms of their ability (noted  $T$  in what follows); as well as in terms of the benefit that they derive from being offered the job (noted  $B$ ). Heterogeneity in  $B$  captures differences in outside options. To fix ideas, it is helpful to think of  $T$  as the score on the Raven test (a reliable predictor of worker performance) and of  $B$  as the monetary net present value of being offered the job (where a negative net present value translates into  $B = 0$ , since being offered the job does not require jobseekers to *take* the job). These are the empirical counterparts that we use for estimation, as described in Section V.

Jobseekers who wish to apply must incur a cost (noted  $C$ ), which is heterogeneous across the population.  $C$  is the net opportunity cost of applying for the job, that is, the economic value of all the things that jobseekers have to give up in order to apply—typically both money and time.<sup>11</sup> This cost is heterogeneous for two reasons. First, the time and money required to make the application differ across jobseekers (e.g. jobseekers who live farther away from the application centre have to pay a more expensive bus fare). Second, the value of time and money differs according to the circumstances of the

---

<sup>11</sup>We also allow  $C$  to be negative. This captures the fact that some people may derive a net benefit from attending the testing sessions, independently of getting the job (e.g. because of the value of networking, or because they learn something valuable about the market).

jobseeker (e.g. poorer jobseekers will find it relatively more expensive to pay the same bus fare compared to jobseekers with better financial resources). If jobseekers dislike being tested,  $C$  will also include the psychic or hassle costs of the application.

We characterise jobseekers along these three dimensions:  $(T, C, B)$ . For simplicity, we assume a finite number  $n$  of ‘benefit types’, so that ability  $T$  and application costs  $C$  have a continuous joint distribution for each of these  $B$ -types. More precisely, we make the following assumptions about the distribution of characteristics across the population of jobseekers.

**Assumption 1** The benefit from receiving a job offer is given by

$$B \in \{b_1, b_2, \dots, b_n\} \text{ where } b_z \geq 0 \text{ for } \{z = 1, 2, \dots, n\}.$$

**Assumption 2** Conditional on  $B = b_z$ , ability  $T$  and application costs  $C$  follow a bivariate normal distribution

$$\begin{pmatrix} T_z \\ C_z \end{pmatrix} \sim \mathcal{N} \left[ \begin{pmatrix} \mu_{T_z} \\ \mu_{C_z} \end{pmatrix}, \begin{pmatrix} \sigma_{T_z}^2 & \rho_z \sigma_{C_z} \sigma_{T_z} \\ \rho_z \sigma_{C_z} \sigma_{T_z} & \sigma_{C_z}^2 \end{pmatrix} \right] \text{ for } \{z = 1, 2, \dots, n\}.$$

Throughout the rest of the paper we use the same notation introduced in Assumption 2. That is, we use sub-indices to denote ability and costs conditional on  $B$ -types. For example,  $\mu_{T_z}$  is mean ability for individuals with benefit type  $z$ . Further, we use  $\sigma_{CT_z}$  to denote the covariance between  $C$  and  $T$ , and  $\rho_z$  to denote the correlation between these two variables.

**Application incentive and wage subsidy.** We model the application incentive as a shock that lowers application costs by an amount  $\tau$ . Similarly, we model the wage subsidy as a shock that raises the value of the job by an amount  $\tau^w$ . In both cases the assumption is that, in line with our empirical findings, the interventions reduce the cost-benefit ratio without affecting jobseekers’ beliefs about the probability of being offered the job upon applying.<sup>12</sup>

**Selectivity and information.** When jobseekers apply for the job, the employer observes their ability  $T$ . Jobseekers, accordingly, make application choices on the understanding that they will get the job if  $T > a$ , where  $a$  captures the perceived selectivity of the recruitment process.

---

<sup>12</sup>We present empirical evidence showing that beliefs about the probability of being offered the job are not affected by treatment in Section IV.E and Table A39. The fact that beliefs do not respond to treatment is consistent with the finding that, in a beauty contest game played during the application process, 80 percent of applicants do not display strategic sophistication (Crawford et al., 2013).

Jobseekers face uncertainty about the likelihood of being offered the job conditional on application. We assume jobseekers know the cost of applying, as well as the benefit of getting the job; but they face uncertainty about recruitment outcomes because they have either imperfect information about the recruitment policy, or imperfect information about their ability. We refer to these two information benchmarks as the ‘noisy selection’ and ‘noisy ability’ cases.

In what follows, we show that in both cases the correlation between cost and ability plays a key role in determining the impacts of application incentives. For a sufficiently large positive correlation, application incentives raise applicant ability regardless of whether the source of uncertainty is ability or selectivity. However, these similar predictions are underpinned by selection decisions made on the basis of different information. In the empirical analysis, we will thus estimate the model under both benchmarks and probe the robustness of our quantitative findings to these different assumptions.

## II.A Analysis

**The noisy-selection case.** To model this case, we assume that ability is known, but selectivity is observed with noise: jobseekers anticipate that the threshold necessary to get the job is a normally distributed random variable with mean  $\mu_a$  and variance  $\sigma_a$ . Thus, an individual with ability  $t$  believes that the probability of being offered the job conditional on an application is  $\Phi((t - \mu_a)/\sigma_a)$ , where  $\Phi(\cdot)$  denotes the standard normal cumulative distribution function.

$B$ -type jobseekers with cost  $C_z = c_z$  and ability  $T_z = t_z$  will apply for the job if and only if

$$(1) \quad \Phi((t_z - \mu_a)/\sigma_a) \geq c_z/b_z$$

From equation 1 it is easy to see that if  $T_z$  and  $C_z$  are not correlated, then an application subsidy that shifts the cost-benefit ratio without affecting beliefs about selectivity, will lead to more but worse applications on average. This is the standard logic that often guides selection policies—a costly application process can improve selection by discouraging the participation of unwanted candidates, who are privately informed about their low ability. Thus, subsidising application costs may reduce average applicant ability.

**Proposition 1.** *Suppose  $(T, B, C)$  are observable and distributed according to Assumptions 1 and 2, with  $\rho_z = 0$ . Further assume jobseekers anticipate that the threshold necessary to get the*

job is  $a \sim \mathcal{N}(\mu_a, \sigma_a)$ . Then it follows that for each  $B = b_z > 0$ , the application incentive (i) increases application rates, and (ii) decreases the average ability of applicants.

*Proof.* Within each  $B$ -type, jobseekers who receive the application incentive face a positive shock that lowers their application costs by an amount  $\tau$ . This means that the application rate among these jobseekers is given by

$$(2) \quad \Pr(C_z \leq \Phi((T_z - \mu_a)/\sigma_a) b_z + \tau)$$

Similarly, the expected ability of applicants in this group is given by

$$(3) \quad E(T_z | T_z \geq \Phi^{-1}((C_z - \tau)/b_z) \sigma_a + \mu_a)$$

If  $C_z$  is not correlated with  $T_z$ , then it is trivial to check that the first expression is increasing in  $\tau$ , and the second expression is decreasing in  $\tau$ .  $\square$

From equations 2 and 3 it is easy to see that application incentives generally operate through two related but different channels. On the one hand, application incentives weakly reduce the ability of marginal applicants at any given level of costs; and on the other hand, they attract additional high-cost applicants at any given level of ability.

When ability and costs are not correlated (i.e. when  $\rho_z = 0$ ) only the first effect matters to determine the impact of the intervention on average applicant ability—this is the standard intuition, captured in Proposition 1. Reducing the cost-benefit ratio increases the number of applications, but lowers the average quality of the applicant pool.

In contrast, when  $\rho_z > 0$ , application incentives activate two forces that act in opposite directions. Lowering the cost-benefit ratio increases applications by attracting a larger share of high-cost applicants, but these new applicants tend to have higher ability on average. If the correlation between ability and costs is strong enough, the second channel dominates and the intervention has a positive selection effect.

In Figure 4 we illustrate this intuition graphically. The application region is shaded in blue ( $(C_z - \tau)/b_z \geq \Phi((T_z - \mu_a)/\sigma_a)$ ), while the superimposed purple contours show the density of the joint distribution of jobseekers ( $\Pr(T_z = t_z, C_z = c_z)$ ). As the figure shows, the application incentive shifts the application threshold to the right, for all levels of ability. However, given the positive correlation between  $T_z$  and  $C_z$ , this expansion of applications adds a larger share of high-cost, high-ability jobseekers. As a consequence, the average ability of applicants increases.

Figure 4 also highlights that, when cost and ability are not correlated, we would expect applicants to have higher average ability than non-applicants. However, the

opposite may be true for a sufficiently strong cost-ability correlation. Indeed, in our experimental data, non-applicants have on average higher ability than applicants (as measured by their GPA score). We will come back to this point in the structural section and show that, while we do not target this moment explicitly, the model nonetheless correctly reproduces it.

< **Figure 4** here. >

**Proposition 2.** *Suppose  $(T, B, C)$  are observable and distributed according to Assumptions 1 and 2. Further assume jobseekers anticipate that the threshold necessary to get the job is  $a \sim \mathcal{N}(\mu_a, \sigma_a)$ . Then it follows that for each  $B = b_z > 0$ , the application incentive (i) increases application rates, and (ii) increases the average ability of applicants, whenever  $(\sigma_{T_z}/(\sigma_a \sigma_{C_z})) \times (b_z/\sqrt{2\pi}) \leq \rho_z$*

*Proof.* See Appendix B

Proposition 2 captures the intuition behind our main experimental result. The application incentive attracts a group of marginal applicants who face larger application costs compared to control group applicants; and because the correlation between costs and ability is positive among jobseekers, these marginal applicants have, on average, higher ability than the applicants in the control group. If the cost-ability correlation is sufficiently large, this indirect channel dominates the standard channel presented in Proposition 1, and the application incentive raises average applicant ability.

This proposition also illustrates the key role played by uncertainty. As noise  $\sigma_a$  grows, the condition on the cost-ability correlation becomes weaker. Intuitively, when the outcome of an application is highly uncertain, application costs play only a limited screening role. In this case, the ‘standard intuition’ effect is small. Thus, a moderate indirect effect is sufficient to obtain a positive impact on applicant ability.

**The noisy-ability case.** Let us now consider the second scenario. Here we assume jobseekers are confident that the selection threshold is fixed at some level  $a$ , but they do not directly observe their ability.<sup>13</sup> Since  $C$  and  $B$  are known, and are potentially informative about ability, we assume that jobseekers can use these variables to update

---

<sup>13</sup>Jobseekers can be wrong about the selection threshold due to overconfidence as in [Spinnewijn \(2015\)](#) and [Banerjee and Sequeira \(2020\)](#). Further, as explained above, we assume that jobseekers do not revise their beliefs about  $a$  in response to treatment. It is possible to show that if we assume instead that  $a$  is the true selectivity threshold, which responds endogenously to treatment, application incentives increase applicant quality if and only if the cost-ability correlation is positive. In Proposition 3 below, we show that a positive cost-ability correlation is similarly required for the case where  $a$  is exogenous. So our key results here are not driven by this assumption.

their belief about the probability of being offered the experiment's job. To fix ideas on this set of assumptions, consider that jobseekers can observe the application conversion rates of people with similar observables, and can use these to make inferences on the probability of a job offer. For example, consider an individual who lives in a central neighbourhood (low  $C$ ) and has a strong educational record (high expected wage/low  $B$ ). Under our assumptions, this person believes that their chances of being offered the experiment's job are the same as those of other applicants from her neighbourhood with a similar educational record.

Formally, we have that  $B$ -type jobseekers with cost  $C_z = c_z$  will apply if and only if  $c_z \leq c_z^*$ , where  $c_z^*$  is the level of costs for which

$$(4) \quad \Pr(T_z > a \mid C_z = c_z^*, B = b_z) = c_z^*/b_z$$

From equation 4 it is clear that in this case the 'standard intuition' that we discussed before is muted. Jobseekers do not directly observe  $T_z$ , so they are not able to self-select directly on ability at the application stage. In fact, if  $C_z$  is not correlated with  $T_z$ , then subsidising applicants would increase the number of applicants, but these new applicants would be a random selection from the pool of jobseekers; so the application incentive would yield no impact on the expected ability of applicants.

Importantly, however, when  $\rho_z > 0$  the indirect channel still operates, because the incentive still attracts on the margin applicants who face larger application costs compared to control group applicants. So in this case too, if costs and ability are positively correlated, the application incentive intervention can yield a positive selection effect. In Figure 5 we illustrate this graphically.

< Figure 5 here. >

The application region is again shaded in blue ( $c_z < c_z^*$ ), and the superimposed purple contours show the density of the joint distribution of jobseekers ( $\Pr(T_z = t_z, C_z = c_z)$ ). As the figure shows, the application incentive shifts the application threshold to the right, for all levels of ability. However, given the positive correlation between  $T_z$  and  $C_z$ , this expansion of applications adds a larger share of high-cost, high-ability jobseekers; and as a consequence, the average ability of applicants increases.

**Proposition 3.** *Suppose  $(T, B, C)$  are distributed according to Assumptions 1 and 2. Assume jobseekers are confident about the selection threshold  $a$ , but they only observe  $C$  and  $B$ , which they can use to update their beliefs about the probability that they will pass the recruitment test  $T > a$ . Then for each  $B = b_z > 0$ , the application incentive (i) increases application rates, and*

(ii) increases the average ability of applicants, if and only if

$$0 < \rho_z < (\sqrt{2\pi}\sqrt{1 - \rho_z^2}\sigma_{C_z})/b_z$$

*Proof.* See Appendix B

Taken together, Propositions 2 and 3 show that the pivotal role of the cost-ability correlation does not depend on whether the uncertainty about the outcome of the application stems from imperfect information over ability, or imperfect information over the recruitment policy. On the other hand, the nature of uncertainty can affect the selection decision in other ways. For example, the ‘standard intuition’ effect only exists in the case where jobseekers know their own ability. Empirically, we do not know which set of assumptions better approximates the information available to jobseekers. Thus, in the structural estimation section, we will bring both information benchmarks to the data and probe the robustness of our findings.

### III The experiment

#### III.A Design

We study the recruitment of workers for clerical jobs in Addis Ababa. The experiment takes place over eight consecutive fortnights. On the Sunday at the beginning of each fortnight, the positions are advertised in a local newspaper and in the main job vacancy boards of the city. The advertisement describes the position as a three-months fixed term appointment based in Addis Ababa and specifies that candidates must hold a university degree or a vocational diploma. Interested individuals are invited to call a specified phone number to get more information about the position and the application process. The deadline for applications is on the Friday of the same week. The positions are based at a local organisation specialised in research and data collection.<sup>14</sup>

A small team of enumerators answers the phone calls of interested jobseekers following a standardised script. First, they ask a short number of questions capturing

---

<sup>14</sup>At the time of the experiment, the organisation employed about 60 permanent workers and 50 fixed-term-contract workers with a similar profile as those recruited for the experiment. All workers hired as part of the experiment were employed for three months and were paid the same wage (which corresponds to the wage offered in the high-wage condition discussed below). This was a surprise for those workers selected under the low-wage offer. The fact that the actual wage would differ from the initial offer was not disclosed to any of the staff involved in the surveys and experiment. The same organisation was contracted for the data collection required for the other parts of this study (e.g. the survey of other employers to elicit beliefs).

callers' socio-demographic characteristics and work experience. Second, they give some information about the position. Third, they explain that, in order to apply for the position the jobseeker has to attend a testing session at a dedicated application centre, on a specified day. Jobseekers have to bring to the session a CV, a cover letter and proof of identity.

We randomly vary two attributes of the position across callers: the wage and whether we offer an application incentive. Callers assigned to the *control group* are informed that the position pays a monthly wage of 1,600 ETB (74 USD) and are not offered the application incentive. Callers assigned to the *application incentive* group are also told that the position pays a wage of 1,600 ETB per month. In addition, these callers are informed that, if they complete the testing session, they will receive a monetary payment of 100 ETB (4.5 USD). This payment is presented as a reimbursement of the costs jobseekers may incur in the application process. Finally, callers assigned to the *high wage* group are told that the position pays a wage of 3,200 ETB (148 USD) per month and are not offered the application incentive. We calibrated these wages at the 35th and 75th percentile of the distribution of earnings for similar positions using data from [Abebe et al. \(2020\)](#). Using jobseekers' assessment of the probability of getting the job, we calculate that the expected value of the high wage offer is worth about 105 USD for the average subject.

All jobseekers who call before the application deadline on a given fortnight are assigned a testing day.<sup>15</sup> This can be from Monday to Friday of the second week of that fortnight, or on the first Monday of the following fortnight. To reduce the risk of contamination across experimental conditions, individuals assigned to different treatment groups are invited to take the test on different days.<sup>16</sup> Each treatment group is assigned to two of these six testing days. This assignment of treatments to testing days is randomly varied every fortnight. If a jobseeker cannot attend the testing session on the proposed day, we allow them to attend the other testing session assigned to his or her treatment group for that fortnight.

We call back all jobseekers four weeks after the first phone call. In this second interview, we ask a set of questions about the job applications that individuals have made

---

<sup>15</sup>We do not allow jobseekers to call on more than one fortnight. After each phone call, enumerators check our database and disqualify the person if they have called on a previous fortnight.

<sup>16</sup>To further reduce the risk of contamination, callers are told that the employer is hiring for multiple positions. This provides a simple explanation for why different callers may be offered different conditions. Specifically, callers in the control group are told that they have been assigned to a position called 'position A'. Callers in the application incentive and high wage groups are informed that they have been assigned to positions 'B' and 'C', respectively. We do not provide any information on why a jobseeker has been assigned to a particular position. If asked about this, the enumerator will reply that one major factor is to have a similar number of candidate for each position.

in the 30 days after the first phone call. Completion of this second phone interview is incentivised with a monetary payment of 20 ETB (.85 USD). We show the full timeline of a typical hiring round in Figure A2 in the Appendix.

We offer three jobs per fortnight – one per treatment group.<sup>17</sup> For each position, the five applicants with the highest score on an index of cognitive ability (which combines the scores on the Raven and Stroop tests) are invited for an interview. The employer decides who among these interviewees is given the job.

As much as possible, the experiment follows the usual recruitment protocols of the employer, which are similar to those of other employers in this market. Importantly, in Addis Ababa it is standard practice not to include information about the wage in the initial advertisement for a position. Further, as explained in Section I, written tests like the one we employ are common.

### III.B Measuring applicant ability

We measure the ability of the individuals that apply for the experiment’s job with a number of popular personnel selection tests. These tests are good predictors of worker productivity and are routinely used by firms worldwide (Chamorro-Premuzic and Furnham, 2010). We also collect information about relevant work experience and GPA scores – two variables that local employers use to screen applicants. Finally, we also collect information about economic preferences.<sup>18</sup>

We measure cognitive ability with the widely used Raven and Stroop tests. The Raven test measures fluid intelligence, the ability to make meaning out of complex information and to reproduce this information. Several meta-analyses have identified the Raven test as the single best predictor of worker productivity (Schmidt and Hunter, 1998; Raven, 2000). This test has been widely used in the recent economics literature to measure worker ability (Dal Bó et al., 2013; Beaman et al., 2013; Abebe et al., 2020). The Stroop test is a popular test of cognitive control, the ability to direct and discipline attention which is required to perform complex tasks (Diamond, 2013). We use a version of the Stroop task developed by Mani et al. (2013).

For non-cognitive skills we use two widely used and validated scales: the big five inventory (BFI-44) and the grit scale (John and Srivastava, 1999; Duckworth et al., 2007). We focus on three facets on non-cognitive ability which have been identified as particu-

---

<sup>17</sup>In a small number of cases, we combine two fortnights of the same treatment group together. When this happens, we offer only one job to the applicants assigned to that treatment group in these two fortnights.

<sup>18</sup>We provide a detailed discussion of each measure in Appendix C.

larly relevant for work performance: conscientiousness, neuroticism and grit. These respectively capture a careful and vigilant attitude at work, the ability to deal with stressful situations, and the capacity to persevere through challenges (Chamorro-Premuzic and Furnham, 2010). We perform standard validity checks for the psychometric measures and satisfy accepted thresholds (e.g. see Table A5 for Cronbach  $\alpha$ ). Laajaj and Macours (2017) emphasise the value of performing validity tests when psychometric scales are used in new contexts. We also administer scales measuring locus of control and confidence.

Further, we collect information about relevant work experience. For this purpose, we use the classification of tasks developed by Autor and Handel (2013). This includes the following categories: physical, routine, problem-solving, managerial, mathematical, and client-interaction tasks. For each of these, we ask participants to report the number of months of experience in jobs that required them to perform that task on a regular basis. We focus on routine, problem-solving and managerial tasks, as these were identified by the employer as the most relevant types of experience.

We aggregate the individual measures in indices of cognitive ability, non-cognitive ability and experience. Each index is constructed as the sum of the standardised values of three measures, as reported in Table A6 in the Appendix (Anderson, 2008).<sup>19</sup>

Finally, we measure four types of economic preferences: an incentivised measure of time preferences, and non-incentivised measures of risk preferences, social preferences and level- $k$  rationality. The task to measure time preferences is an adapted version of the game by Augenblick et al. (2015). In this task, participants have to allocate units of work across two work sessions. For risk preferences and social preferences we use questions from the Global Preferences Survey (Falk et al., 2016b). Finally, we administer a simplified and non-incentivised version of the beauty contest game to elicit level- $k$  rationality (Crawford et al., 2013).

### III.C The sample, randomisation and attrition

Over the eight fortnights of the experiment, 4,689 jobseekers called to inquire about the position — an average of about 590 individuals every fortnight. The number of fortnightly callers stayed constant over the course of the experiment, suggesting that

---

<sup>19</sup>We think of the three components of the index as representing three distinct facets of a particular ability. We thus give each component of the index equal weight. Results, however, are qualitatively unchanged if we weight by the inverse of the covariance matrix. Further, ideally, we would standardise each test score by the variance of that test for a pool of applicants attracted at the average wage. We are unable to do this, as we either attract applicants at a below-average or at an above-average wage. We thus use the variance of a given test among all applicants in order to standardise the value of that test.

the positions generated sustained interest among jobseekers. Table 1 reports summary statistics for the population of individuals who called to inquire about the position. The typical caller is young, male and has some work experience. The average age is 26. 15 percent of the sample is 30 or older. Women account for 21 percent of the sample. On average, callers have 28 months of wage-work experience. This masks substantial heterogeneity, as 47 percent of the sample has no work experience. Callers also have a variety of educational backgrounds.

< Table 1 here. >

We randomise the offer of the two treatments using a stratification rule in order to improve covariate balance (Bruhn and McKenzie, 2009). We create strata of six consecutive callers of the same gender and same level of work experience.<sup>20</sup> In each stratum, we randomly allocate two callers to the control group, two callers to the application incentive group and two callers to the high wage group. These callers are invited to a testing session at the application centre during the following week. There are two testing sessions per treatment group, per fortnight. We randomise the allocation of testing sessions to days of the week. We do this in a single draw for all eight fortnights and re-randomise until we have an allocation that is balanced across days of the week.<sup>21</sup>

We find that covariates are balanced across treatment groups and that attrition is modest and uncorrelated with treatment. 1,557 callers are assigned to the control condition, 1,559 to the incentive condition, and 1,573 to the high wage condition. Table 1 reports descriptive statistics and balance tests for the characteristics of the callers that we measure during the first phone interview. The joint orthogonality tests show that, overall, characteristics are balanced across treatment groups. Individual covariates are also generally balanced.<sup>22</sup> Further, appendix Table A9 shows that the assignment of individuals to treatment is strongly balanced across weeks. In terms of attrition, in the second phone survey, we interview 93.5 percent of the sample (attrition is thus 6.5 percent). This is consistent with recent studies with similar populations in urban East Africa (Abebe et al., 2020). Figure A3 shows that attrition is not systematically related

---

<sup>20</sup>We define an experience dummy using the median number of months of work experience of callers in the pilot.

<sup>21</sup>The randomisation rule is that (i) each treatment should be allocated two testing days each fortnight, and (ii) no treatment should be allocated, overall, more than three or less than two sessions on the same day of the week. For this exercise, we consider the Monday session on the the following fortnight as being a distinct 'day of the week'.

<sup>22</sup>The only exception is a 3 percentage point difference between the high wage group and the control group in the proportions of callers who are unemployed and wage employed.

to treatment status and Table A8 confirms that the sample remains balanced after attrition.

Using data from the 2013 Labour Force Survey (LFS), we also show that our experimental sample resembles in many ways the population of jobseekers in Addis Ababa. In particular, average age and work experience in our sample match those of a comparable sample of jobseekers from the LFS, as we show in Table A10 in the online Appendix. We find some differences in terms of gender (women are under-represented in our sample) and unemployment duration (the long-term unemployed are under-represented in our sample). These differences are likely due to the fact that women and the long-term unemployed search at lower intensity and hence are less likely to respond to the ad placed by the employer. As we will show in Section IV, these groups of jobseekers tend to drive the treatment effects on applicant quality. Our estimates are thus likely to be a lower bound of the treatment effects that would be observed in a more representative sample.

### III.D Empirical strategy

Our objective is to study the impacts of the interventions on application rates and on the ability of applicants.<sup>23</sup> We estimate effects on application rates using a regression model of the following form:

$$(5) \quad \text{apply}_i = \beta_0 + \beta_1 \cdot \text{incentive}_i + \beta_2 \cdot \text{high wage}_i + \kappa_b + u_i,$$

where  $\text{apply}_i$  is a dummy that captures whether person  $i$  has applied for the job,  $\text{incentive}_i$  and  $\text{high wage}_i$  identify individuals who have been offered the application incentive and the high-wage treatment, and  $\kappa_b$  are stratum dummies (Bruhn and McKenzie, 2009). The coefficients  $\beta_1$  and  $\beta_2$  capture the change in application rates generated by the application incentive and the high wage offer. This model is estimated over the *baseline sample*, that is, over the sample of individuals who called to inquire about the position. We use a similar model to study the effects of the interventions on expectations and other job-search activities. These outcomes were measured during the second phone call and thus we use the second phone call sample to run this additional analysis.<sup>24</sup>

We study impacts on the ability of applicants by measuring changes in the conditional mean and conditional quantiles of the measures of quality discussed in the previous section. For this analysis, we restrict the sample to all applicants. We measure

<sup>23</sup>All reduced-form analysis was pre-registered. In the Online Appendix, we list a number of deviations from our original plan motivated by data and fieldwork challenges.

<sup>24</sup>As discussed in the previous section, attrition between the two surveys is limited and uncorrelated with treatment. Furthermore, the sample remains balanced after attrition.

changes in mean applicant ability by estimating a simple OLS regression model of this form:

$$(6) \quad y_i = \delta_0 + \delta_1 \cdot \text{incentive}_i + \delta_2 \cdot \text{high wage}_i + u_i,$$

where  $y$  is a measure of worker ability. Further, we estimate a conditional quantile function of the following form, using a standard quantile regression model (Koenker and Hallock, 2001):<sup>25</sup>

$$(7) \quad Q_\theta(y_i|X_i) = \gamma_0 + \gamma_1 \cdot \text{incentive}_i + \gamma_2 \cdot \text{high wage}_i.$$

For each measure of worker quality  $y$ ,  $\gamma_1$  and  $\gamma_2$  capture the change in conditional quantile  $\theta$  caused by the treatments. For example, suppose that we are studying the 90th percentile of the distribution of cognitive ability and that we obtain an estimate of  $\gamma_1$  of 1. This would say that an applicant at the 90th percentile of the distribution in the incentive group has a cognitive ability score that is one point higher than an applicant at the 90th percentile of the control distribution. A key implication of this quantile shift is that the proportion of applicants who score above the 90th percentile of the control distribution increases. This suggests that, to study changes in applicant ability, we can also compare the probability that an applicant scores above a given threshold across the two applicant groups. In the results section, we show that our findings are robust to the use of this alternative empirical strategy.

We focus the quantile analysis on five percentiles: 90th, 75th, 50th, 25th, 10th. We also present a test of stochastic dominance first proposed by Barrett and Donald (2003). Stochastic dominance occurs when the CDF of one distribution is weakly smaller than the CDF of the other distribution at all points of the support *and* strictly smaller at least at one point. The null hypothesis of the Barrett and Donald (2003) test is that the CDF of one distribution is weakly smaller than the CDF of the other distribution. To have evidence that distribution A dominates distribution B, we should thus both (i) reject that B is weakly smaller than A and (ii) fail to reject that A is weakly smaller than B. In the results section, we report and interpret the findings of both tests.

We report robust standard errors throughout the paper.<sup>26</sup> Further, we control for multiple-hypothesis testing by using sharpened  $q$ -values that control the false discovery rate (Benjamini et al., 2006).<sup>27</sup>

<sup>25</sup>In models (6) and (7) we do not include the randomization block dummies. This is because we estimate these models on the sample of applicants, which, by design, does not include all of the individuals that were originally randomised into the three experimental conditions.

<sup>26</sup>For quantile regressions, robust standard errors are computed using the methods proposed by Machado and Silva (2000) and the Stata command developed by Machado et al. (2011).

<sup>27</sup>We have three indices of ability, so each individual test is repeated three times, forming a family

## IV Results

### IV.A Impacts on application rates

We find that the incentive has a large and significant effect on applications. Individuals in the incentive group are 11.5 percentage points more likely to apply for the position than individuals in the control group. 41 percent of subjects in the control group apply for the position, so this treatment effect amounts to a 28 percent increase in application rates. Further, we find that individuals in the high-wage group are 18.6 percentage points more likely to apply to the position. Thus the application incentive generates an increase in applications that is about two thirds of the increase in applications that can be obtained by doubling the wage. The two effects are statistically different from each other. We report these results in Table 2.

< Table 2 here. >

As shown in Figure A4, impacts on application rates are stronger in the upper part of the GPA distribution. For example, the increase in applications generated by the incentive treatment in the lowest decile of the GPA distribution is close to 8 percentage points. In the top decile of the distribution, the increase is a significant 14.5 percentage points. These observations suggest that the overall effect on the quality of the pool of applicants, both at the mean and at the top of the distribution, is positive. We discuss these results in detail in the next subsection.

### IV.B Impacts on the quality of the applicant pool

The application incentive improves the quality of the applicant pool. This is our most important finding. The incentive raises average cognitive ability among applicants by .25 points, or .12 of a standard deviation (Table 3). This effect is significant at the 5 percent level and is robust to the correction for multiple comparisons.<sup>28</sup> Applicants in the incentive group perform significantly better in both the Raven and the Stroop tests. Compared to applicants in the control group, they answer correctly 1.2 additional questions in the Raven test and they require 2.6 fewer seconds to complete the Stroop of hypotheses. To calculate  $q$ -values, we first compute standard  $p$ -values for each test in a given family. Then, we run the sharpened procedure proposed by Benjamini et al. (2006) using these  $p$ -values. The  $q$ -values we obtain express the expected proportion of false discoveries that we need tolerate if we want to reject the null hypothesis of a given test.

<sup>28</sup>This estimate is also not sensitive to the exclusion of tests carried out on specific days of the week. We show this in Figure A6 in the Appendix. Also, the findings in this section are qualitatively unchanged when we weight the index by the inverse of the covariance matrix (Table A17, A18, A19).

task. These treatment effects are of a broadly similar magnitude as those documented in previous worker selection experiments. For example, Dal Bó et al. (2013) document an increase in performance on the Raven test of about half a correct answer. We report the full results for the individual tests in Table A11 in the Appendix. We also find that the applicants attracted by the incentive have GPA scores that are a significant .1 standard deviation higher than control applicants (Table A12). This is an important result as many firms in Addis Ababa use GPA scores to assess candidates' ability. Thus the applicant pool improves also in terms of the screening criteria used by local firms.

The increase in quality occurs both at the top and at the bottom of the distribution. These results hold when looking at quantile shifts, number of top and bottom applicants, and average quality of the top applicants. We look at each of these in turn. First, we find that the cognitive ability scores at the 90th, 75th and 25th percentiles improve significantly (Table 3). These effects correspond to about .1 a standard deviation of the cognitive ability index; the respective  $q$ -values are always below .15. We also estimate positive, but insignificant effects at the 50th and 10th percentiles. Overall, a Wilcoxon rank-sum test rejects the equality of the cognitive ability distribution in the control and incentive groups ( $p=.038$ ). Second, we document a large effect on the number of top applicants (defined as individuals above the 90th percentile of the cognitive ability score in the control group). Top applicants nearly double from 63 in the control group to 117 in the incentive group. This effect is generated by a combination of higher application rates, and a significant, 5.3 percentage points increase in the proportion of top applicants in the applicant pool (see Table A13 in the Appendix). At the same time, the number of applicants at the bottom of the distribution is fairly stable. For example, compared to the control condition, the application incentive attracts only nine additional applicants who score below the 10th percentile of the control distribution. Third, we find that the average ability of top applicants increases significantly. In Table A14 we show regressions of average cognitive ability and GPA scores for samples comprising the top 20, 10 and 5 applicants for each job (the top 5 applicants are invited for the interview). We document sizeable increases in both scores between .2 and .5 standard deviations. In particular, the application incentive increases the GPA scores of the top 5 applicants by a significant .5 standard deviations and the cognitive ability score by a large but insignificant .3 standard deviations.

Consistently with the results for specific quantiles, we find suggestive evidence that the cognitive ability distribution among treated applicants stochastically dominates the control distribution. This is an attractive feature if the firm's objective is to maximise

the ability of its hires.<sup>29</sup> We see the characteristic pattern of stochastic dominance when we plot the cumulative distributions of cognitive ability for the two groups (Figure 6). We find a similar pattern if we look at the results of the Raven test, the Stroop test and GPA separately (Figures A8, A9 and A10 in the Online Appendix). Using the formal test of Barrett and Donald (2003) we find no evidence to reject the hypothesis that the CDF of the incentive distribution is weakly smaller than the CDF of the control distribution ( $p=.949$ ). This result is consistent with dominance of the incentive distribution over the control distribution. However, it also consistent with the equality of the two distributions. We thus also test the null hypothesis that the CDF of the control distribution is weakly lower than the CDF of the incentive group. For this test we obtain a  $p$ -value of .136, giving us suggestive evidence of stochastic dominance.

The high-wage offer also attracts an applicant pool with higher cognitive ability. We estimate significant positive effects at the mean, and at the 90th, 75th and 25th percentiles (see Appendix Figure A11 for the full distribution). The magnitude of these point estimates is smaller than those we obtained for the application incentive, but we cannot reject the null hypothesis that the two treatments have the same effect. The significant estimates of the impact of the high-wage offer are associated with  $q$ -values above .1 (and in two cases above .2), suggesting these results are generally not robust to the multiple-comparisons correction.

< Table 3 here. >

< Figure 6 here. >

Lastly, we are unable to find significant differences in non-cognitive ability or experience between applicants in the incentive group and applicants in the control group. The high-wage offer significantly increases median non-cognitive ability, but does not significantly affect the other percentiles of the distribution. Tables A15 and A16 report the results from these regressions.<sup>30</sup>

---

<sup>29</sup> Stochastic dominance makes it possible to unambiguously rank distributions for objective functions that are increasing in the value of the random variable (Deaton, 1997; Barrett and Donald, 2003). Thus, in our setting, the dominant distribution would be preferred both by firms who maximise the expected quality of hires, and by 'risk-averse' firms with an objective function that is increasing and concave in quality. The comparison would not be unambiguous, however, if firms value having a smaller pool of applicants or if acceptance rates are lower in the dominant group. We consider the first point in Section V. Regarding the second point, we show below that the increase in quality generated by the incentive is concentrated among those jobseekers with the weakest outside options. These jobseekers are likely to have the highest acceptance rates. This further increases the value of the applicant pool attracted by the application incentive.

<sup>30</sup>We are similarly unable to document systematic impacts on ability when we ipsatise the psychome-

## IV.C Impacts on search for other jobs and job-search outcomes

We do not find evidence that the application incentive distorts individuals' search for other jobs or impacts their labour market outcomes. This is not surprising, as the small cash incentive ensures that applying to the experiment's job does not deplete the resources available to search for other positions. To study the search for other jobs, we use the data collected during the second phone interview, 30 days after the initial phone call, and a regression model with same form as model (5). We investigate whether the interventions change the number of applications made, the amount of money and time spent on job search, the number of interviews and job offers obtained, and whether the jobseeker is currently working in a new job. We report these results in Table A21 in the Appendix.<sup>31</sup> For the application incentive, we consistently estimate small and insignificant coefficients.

On the other hand, we find that individuals in the high-wage group have significantly worse outcomes than the controls: they obtain .04 fewer interviews, .03 fewer offers and are about 2 percentage points less likely to be working in a new job. In Table A23 we also document that individuals in the high-wage intervention are significantly less likely to have obtained a job in their desired occupation, a job that they see themselves doing in the long run or a job with an open-ended contract. One possible explanation for these results is that the individuals that are induced to apply to the experiment's job by the higher wage run out of resources to search for other jobs. The effects of the high-wage offer on search effort are indeed negative: the intervention is associated with a significant 10 percent decrease in the number of visits to the job vacancy board — a costly form of search that requires the use of public transport (Table A24). In Table A21 we also document effects on applications that point in the same direction, but are smaller and statistically insignificant: a 4 percent decline in the number of applications to other jobs, and a 3 percent decline in the time spent on job applications.

---

tric measures to correct for acquiescence bias (the tendency to agree with any statement, regardless of the content of that statement). We report these additional regressions in Table A20.

<sup>31</sup>In this Table, we exclude the application to the experiment's job in the definition of the dependent variables. In Table A22, on the other hand, we report estimates for the *total* number of applications, which includes the application to the experiment's job. Under this alternative definition, we estimate that both interventions are associated with a marginally insignificant increase of about .1 applications per jobseeker.

#### IV.D Mechanisms: Who drives the increase in quality?

We study the heterogeneity of treatment effects along several dimensions. These include demographic characteristics (gender and age), labour market variables (employment status and work experience) and a variable capturing the net present value of the job to each jobseeker. We calibrate this variable by comparing the control wage of the experiment’s job to a forecast of the wage that each individual can expect to be paid in the market, based on their observables. We describe in detail the calibration procedure in Appendix F. For each dimension of heterogeneity  $x$ , we estimate a model of this form:

$$(8) \quad \begin{aligned} y_i = & \beta_0 + \beta_1 \cdot \text{incentive}_i \cdot I(x_i = 1) + \beta_2 \cdot \text{high wage}_i \cdot I(x_i = 1) \\ & + \beta_3 \cdot \text{incentive}_i \cdot I(x_i = 0) + \beta_4 \cdot \text{high wage}_i \cdot I(x_i = 0) \\ & + I(x_i = 1) + \kappa_b + u_i. \end{aligned}$$

Model (8) gives us separate estimates of the effect of treatment for individuals for whom  $x = 1$  and individuals for whom  $x = 0$ . When a variable is continuous,  $x$  is dummy that splits the sample at the median of that variable. For each regression and each treatment, we present a test of the hypothesis that there is no heterogeneity in the effect of that treatment ( $H_0 : \beta_1 = \beta_3$  for the incentive, and  $H_0 : \beta_2 = \beta_4$  for the high-wage offer). We bootstrap the standard errors of the regressions studying heterogeneity with respect to the value of the job, to reflect the fact that we are using a generated regressor. Results are reported in Tables A25 to A33 in the Appendix.

We find that the increase in cognitive ability caused by the incentive is significantly stronger among women, the unemployed, the less experienced, and for those individuals whom we estimate to value the job the most (Table A27). Impacts are also much larger for applicants below the median age, but in this case we cannot reject a null hypothesis of no heterogeneity ( $p = 0.162$ ). These are all groups that on average fare worse in the labour market and that respond more strongly to job-search support (Card et al., 2010; Abebe et al., 2020). Further, with the exception of work experience, we cannot document heterogenous impacts of the high-wage offer with respect to these dimensions. The magnitude of the differences in impacts across groups is often large. For example, among males, the effect of the incentive on average cognitive ability is close to zero. Among women, on the other hand, the cognitive ability score more than doubles (and the Raven test score increases by about 4.7 points or about .4 of a standard deviation, as reported in Table A28). We also document significantly larger effects for women at the 90th and 75th percentiles of the cognitive ability distribution (Table A31). We illustrate these results graphically in Figure A7, where we show that the proportion of female top applicants grows from 21 percent in the control group to 33 percent in the application

incentive group.

On the other hand, we do not find systematic heterogeneity in the treatment effect on the probability of applying for the position (Table A25). For each dimension of heterogeneity, we can use the estimated impacts on applications rates and group-specific quality to break-down the total treatment effect in three parts: a compositional effect, a within-group effect for the first group, and a within-group effect for the second group. This exercise, which we report in Table A32, shows that the impact of the intervention is largely driven by within-group improvements in quality among those groups that have weaker employment prospects (women, youth, the unemployed, the inexperienced). This suggests that the effectiveness of the intervention could be improved by targeting the subsidy towards these groups, a point which we explore in depth in Section V. Further, when we use all observables to obtain a forecast of market wages and then calculate the present value of the job, we find that virtually all of the increase in ability is driven by individuals with a low forecasted market wage and hence a high present value of the job. This finding helps us rule out a mechanism whereby the incentive increases ability by attracting individuals with strong outside options and hence a relatively modest personal gain from getting the job. On the contrary, it suggests that the application subsidy attracts high-ability individuals with weak outside options and thus has the potential to increase allocative efficiency.

Table A32 also shows that compositional effects play a minor role in explaining our results. In Section V, we show that our model can reproduce this feature of the experimental findings. The absence of compositional effects is explained by the fact that costs are highly heterogeneous within demographic groups. This means that there is substantial overlap between group-specific cost distributions and, thus, that in all demographic groups there is a sizeable number of individuals who face high application costs.

## IV.E Alternative explanations

In this section we consider seven alternative explanations for our results that are unrelated to application costs. We do not find evidence suggesting that these channels drive the effect of the application incentive.

**Do the interventions change test effort?** First, we consider whether the treatments affect test effort. This could happen for a number of reasons. For example, applicants in the wage intervention may exert more effort in the selection test since the higher salary makes the position more valuable. We study this alternative hypothesis in two ways.

First, we note that the interventions have significant treatment effects on GPA (reported in Table A12), a measure of ability that was determined before treatment and that hence cannot be confounded by differential test effort. Second, we follow Dal Bó et al. (2013) and collect a direct measure of test effort through a task that requires effort, but very little ability. In this task, applicants have to transcribe ten strings of meaningless letters. In Table A36 we study whether two measures of performance on this task — the number of strings that have not been transcribed and the number strings that have been transcribed incorrectly — improve with treatment. We find that the average applicant in the control group fails to transcribe 0.08 out of 10 strings and transcribes incorrectly 0.7 out of 10 strings. Applicants in both treatment groups perform in a similar and statistically indistinguishable way to control applicants, both in terms of mistakes and in terms of number of strings transcribed. For example, the average applicant from the high-wage intervention makes an additional .067 mistakes — a very small difference. We obtain similar results if we study the likelihood of making any mistake or failing to transcribe any string. Altogether, this evidence suggests that the treatments do not change test effort.

**Do individuals in the incentive group apply primarily to collect the monetary payment?** Next, we explore whether application incentives attract individuals that are primarily interested in the immediate monetary payment, and not in the position. We consider two implications of this hypothesis. First, if this hypothesis was correct, we would expect lower test effort among treated applicants. However, as discussed above, applicant test effort is very similar across experimental groups. In Figure A12 we further show that the distribution of test effort in the control and incentive groups looks remarkably similar at all point of the support (and a Kolmogorov-Smirnov test fails to reject the equality of these distributions). Second, we would expect that, compared to the high-wage condition, incentives would generate an increase in application rates that is more skewed towards low GPA jobseekers, who have a low probability of securing the position and should hence discount the wage increase more heavily. This is the opposite of what we observe when we plot application rates against GPA (Figure A5).

**Does the application incentive help present-biased individuals commit?** Third, we explore the hypothesis that application incentives help present-biased individuals follow through on their intention of applying to the experiment's job. On the day the application is due, a present-biased individual who plans to take the test may be tempted to deviate from his or her plan. The incentive makes it more costly for them to change their mind — as they would forgo both the future opportunity of getting the job *and* the

immediate monetary payment — and can thus help them act in a more time-consistent manner. While plausible, this explanation is not supported by the data. We do not find evidence suggesting that the incentive attracts applicants who have different time preferences, which we measure through an incentivised task similar to the task proposed by [Augenblick et al. \(2015\)](#) (described in detail in Appendix C). In Table A34, we report structural estimates of *average* present bias, discounting and cost of effort, which are obtained by pooling all individual decisions in the incentivised task. We find that, on average, individuals have a value of the present-bias parameter  $\beta$  that is less than one, indicating time consistency. This is very similar across treatment groups. In the same table, we further show that the discounting and cost-of-effort parameters of applicants from the incentive group are similar and statistically indistinguishable from the controls. The high-wage treatment, on the other hand, attracts significantly more patient applicants. This is not surprising, as the present value of the high-wage offer is greater for patient individuals. In Table A35, we then present additional regression results for *individual* measures of present bias, obtained by running a separate estimation for each applicant. About 30 percent of applicants are present biased according to this measure. This share goes up by an insignificant 2 percentage points among treated applicants. Thus, overall, the evidence does not support the hypothesis that the effect of the incentive is driven by present biased individuals. In Table A35, we also show that there are no significant differences in other economic preferences (e.g. risk and social preferences), which we measure through non-incentivised questions.

**Do the interventions make the job more salient?** Fourth, we study whether the treatments induce jobseekers to pay more attention to the experiment’s job. In our context, individuals may suffer from stress and cognitive load. As a result, they may be inattentive, especially when the benefits of an action accrue in the future, and may thus forget to apply to some of the jobs they are interested in ([Mullainathan and Shafir, 2013](#)). The interventions could induce jobseekers to pay more attention to the experiment’s job because they increase the cost of forgetting, or because they make the job stand out among others ([Gabaix and Laibson, 2005](#); [DellaVigna, 2009](#)). This increased salience would reduce the likelihood that individuals forget to apply to the position, potentially driving some of the impacts on application rates that we have documented.

We study this alternative explanation in two ways. First, we note that a mechanism of this type is likely to work against the direction of our findings on ability, as cognitive load temporarily decreases cognitive ability ([Mani et al., 2013](#)). Second, we directly test this explanation by leveraging the fact that salient information is more likely to be

remembered (Botta et al., 2010; Santangelo and Macaluso, 2013). If the interventions make the position more salient, we would expect treated jobseekers to remember more accurately the information about the job that was given to them. We implement this test in Table A37, using a question in the second phone call that asks jobseekers to recall the wage that was discussed in the first phone interview. In the control group, we find that about 69 percent of individuals report the correct wage. The remaining subjects either report an incorrect figure, or declare that they do not remember. The average report has an absolute mistake of 167 ETB. Importantly, we are unable to find systematic evidence of better recollections for the application incentive group. However, we find that individuals in the high-wage group recollect the wage of the position more accurately: they are more likely to report the correct figure (by 3.9 percentage points), and they make smaller absolute mistakes (by 45 ETB on average). These results are robust to controlling for whether individuals have applied for the job.<sup>32</sup> Thus, overall, while these results suggest that the wage intervention increases salience, the evidence does not support the hypothesis that this is also the case for the incentive intervention.

**Do the interventions change jobseekers' beliefs about their prospects in the labour market?** We study whether individuals update their beliefs about their prospects in the labour market. This could be the result of a revision in the beliefs that individuals hold about their own employability, or in the beliefs about the state of the labour market. For this test we use two questions from the second phone interview. In the first question, we ask subjects to forecast the number of weeks that it would take them to find a job that paid at least their reservation wage. In the second question, we ask respondents to report the wage that they expect this job will pay.<sup>33</sup> We find that the application

---

<sup>32</sup>Applicants have a second chance to inquire about the wage during the application process and may thus have better recollections for reasons unrelated to salience. Since the treatments increase application rates, this may bias upwards the coefficients in columns (1) and (2) of Table A37. To correct for this, the models reported in columns (3) and (4) control for whether the individual has applied for the job. This control variable captures an endogenous mediator and hence the models in columns (3) and (4) should be interpreted as a basic form of mediation analysis (Acharya et al., 2016). We present a formal mediation analysis at the end of this section.

<sup>33</sup> To elicit expectations about the wage, we follow the method of Attanasio and Kaufmann (2009). We ask respondents to report the minimum and maximum wage that the job can pay. We then identify the midpoint between these two values and ask respondents to report the probability that the job will pay more than the midpoint. Following Attanasio and Kaufmann (2009), we assume that beliefs follow a triangular distribution. This distribution is fully characterised by an upper bound, a lower bound and a mode. The maximum and minimum wage reported by respondents identify the upper and lower bounds. Given the two bounds, the value of the CDF at the midpoint identifies the mode of the distribution.

incentive does not have a significant effect on either of these beliefs. The high-wage offer, on the other hand, significantly increases expected wages by about 8.5 percent. Table A38 in the Appendix reports these results.

We also do not find evidence that the interventions change jobseekers' beliefs about the *probability of getting the experiment's job*. We show this result in Table A39 in the Appendix. This finding is consistent with the low levels of strategic sophistication that are documented in a simplified beauty contest task (Crawford et al., 2013). In general, applicants are overconfident about the likelihood of getting the experiment's job. This is consistent with recent research showing that beliefs about individual performance are characterised by overconfidence in several contexts, including job search (Malmendier and Tate, 2015; Spinnewijn, 2015; Banerjee and Sequeira, 2020).

The impacts on beliefs do not vary significantly according to the ability of the job-seeker (as proxied by their GPA score) or previous work experience. We show this in Figures A13 and A14 in the Online Appendix, where we report impact estimates for various subgroups and a  $p$ -value for a test of the null hypothesis of no heterogeneity.

**Do the interventions change jobseekers' beliefs about the attributes of the job?** We test whether the treatments affect the beliefs that individuals hold about the characteristics of the experiment's job. To test for this, in the second phone call we collect jobseekers' beliefs about several attributes of the job: holidays, non-standard working hours, the degree of autonomy, how satisfying the work will be, whether they will learn new skills, etc... We regress each of these beliefs on the two treatment dummies and report results in Table A40 in the Appendix. We find that the application incentive has a modest significant effect on two of these dimensions: the proportion of people who think the job will have more than four days of holidays per month goes up by 2.5 percentage points, and the proportion of people who think that the job will help them to find a job in the future goes up by 3.4 percentage points. These two expectations are weak predictors of the decision to apply for the experiment's job. Among control group individuals, the belief that the job has long holidays raises the probability of making an application by 7.8 percentage points, while the belief that this job will help with job search in the future raises the probability of making an application by 8.2 percentage points. To assess the potential effect of this channel on application rates, we multiply the treatment effects on the beliefs by the effects that these beliefs have on application rates and add up. The result is that this channel can explain a change in application rates of about half a percentage point. In other words, net of this effect, the application incentives would raise applications by 11 percentage points (as opposed to 11.5

percentage points).

We provide additional evidence that application incentives do not change beliefs about job attributes with a survey experiment on a new sample of jobseekers. We select 724 jobseekers through random visits at the same job-vacancy boards where the experiment's job was originally advertised.<sup>34</sup> During the survey, we give subjects some information about a hypothetical job, including salary and nature of the job, randomising whether this description mentions a monetary application incentive or not. We then ask respondents whether they expect the job to have a number of attributes (since the questions refer to a hypothetical job, respondents are not rewarded for correct answers). In Table A42 we show that, in this new sample, application incentives do not generate any significant or sizeable change in job-attribute beliefs. Together with the results on retrospective beliefs reported in Table A40, this evidence points against the possibility that the impacts of the application incentive intervention are driven by a misperception of job attributes.<sup>35</sup>

**Mediation analysis.** We use mediation analysis to quantify the contribution of the changes in beliefs and salience caused by the wage intervention to its overall impact on application rates. We do this by estimating the *average controlled direct effect* (ACDE) of this intervention (Vansteelandt, 2009; Acharya et al., 2016). This quantity is defined in a framework where a treatment can have both a direct impact on the outcome of interest and an indirect impact which runs through a mediator. For example, the wage intervention may have a direct impact on application rates as it makes the position more lucrative, but also an indirect effect if it makes potential applicants more confident. The ACDE is defined as the direct effect of the intervention, that is, the effect that the

---

<sup>34</sup>We carried out this survey between December 2019 and January 2020, as part of a different project. We report basic descriptives for this new sample in Table A41. The jobseekers in this sample have broadly similar characteristics to the jobseekers in our original experiment: they are 24 years old on average, 20 percent of them were born in the capital city, about 40 percent have work experience and, those who do, have worked for 20 months on average. The corresponding figures for our original sample are: 26 years of age, 24 percent born in Addis Ababa, 53 percent have work experience and, those who do, have worked for 28 months on average. The two samples however differ in terms of gender: 49 percent of respondents in the new sample are female, as the sample was stratified by gender, while only 21 percent of the respondents in the original experiment are female. Table A41 also reports a standard battery of balance tests. The overall test of orthogonality confirms that treatment assignment is balanced. However, we find some imbalance ( $p=0.08$ ) in terms of the proportion of jobseekers who are currently unemployed.

<sup>35</sup>As an additional check, we debriefed the staff from the original experiment to find out whether treated applicants asked more or different questions about the position during the first phone call. The staff was unable to recollect any systematic differences in the questions asked by treated and control subjects.

intervention would have if the mediator was not allowed to respond to treatment and hence the indirect effect was removed.<sup>36</sup> Here, we focus on the ACDE on application rates and consider a list of potential mediators that we showed to be affected by the interventions in the analysis above: beliefs about labour market prospects, beliefs about the attributes of the experiment’s job, salience of the position.

We find that the ACDE of the high-wage offer on application rates is 9 percentage points (with a 95% confidence interval ranging from 3 to 13 percentage points). This is significantly smaller than the original estimate reported in Table 2 (which was 18.7 percentage points, with a 95% confidence interval ranging from 15 to 22 percentage points), suggesting that the mediators have a quantitatively large influence on application rates in the high-wage group. The controlled direct effect of the incentive, on the other hand, is quantitatively similar and statistically indistinguishable from the original treatment effect (the demediated effect is 13 percentage points, slightly larger than the original effect of 11.5 points). This is not surprising, as we only find evidence of large and systematic effects on the mediators for the wage treatment.

#### IV.F Other characteristics of applicants

We report one final set of results on the characteristics of applicants. We are particularly interested to shed light on a set of productivity-relevant traits that may not be captured by our main indicators of ability. These include confidence and motivation, and personal constraints that may affect the ability to work. In Table A43, we report regressions for a comprehensive battery of psychometric measures — grit, locus of control, core self-evaluation, self-esteem as well as the individual big-five traits. We are unable to find any meaningful differences between control and incentive applicants: all effect sizes are markedly below .1 standard deviation and are precisely estimated. Further, in Table A34 and Table A35 we show that treated applicants have similar economic preferences and a similar cost of effort as control applicants. As explained above, our measure of cost of effort is obtained in a real-effort task where individuals have to allo-

---

<sup>36</sup>In order to identify the ACDE we have to assume *sequential unconfoundedness*. In a case where treatment is randomly assigned, this amounts to assuming that there are no omitted variables which confound the effect of the mediator on the outcome, conditional on treatment and a set of pre-treatment controls (Acharya et al., 2016). Given this assumption, we can identify the ACDE with a simple two-step procedure. In the first step, we regress the outcome on the mediator, the treatment dummies, a set of controls, and the interaction between the mediator and all other variables. We then obtain the predicted value of the outcome fixing all mediators to zero. This is the ‘demediated’ outcome. In the second step, we regress the demediated outcome on the treatment dummies. The coefficients from this regression give us the estimate of the ACDE.

cate a certain amount of work across two different work sessions. In this task, treated applicants make similar allocation decisions as control applicants — a result that suggests that they have a similar ability to schedule work in advance and to work without interruption as control applicants. Overall, the results of this further analysis suggest that, apart from the differences in cognitive ability highlighted above, the interventions attract pools of applicants that have characteristics similar to those of control applicants.

## V Structural analysis

In this section we discuss the identification and estimation of a structural model, based on the theoretical framework presented in section II. Our estimated structural parameters indicate that application costs are large, heterogeneous and positively related with ability. These results are consistent across a number of robustness tests.

### V.A Identification and estimation

Our objective is to estimate the parameters that characterise the joint distribution of costs ( $C$ ) and ability ( $T$ ), for each level of the value of the job ( $B$ ). In our main empirical model, we allow for two possible values of  $B$ . Thus, we need to identify ten parameters for the joint distribution of  $T$  and  $C$  of low and high- $B$  jobseekers. We also estimate the magnitude of the shocks to costs and benefits of the two interventions ( $\tau$  and  $\tau^w$ ) and the parameters that capture perceived selectivity. In the case where ability is noisy, we estimate a single selectivity parameter  $a$  (which may differ from true selectivity if jobseekers have an inaccurate perception of the selection process). In the noisy selection case, we assume that jobseekers believe selectivity is an independent, normally distributed variable with mean  $\mu_a$  and standard deviation  $\sigma_a$ . Thus, we have a total of thirteen parameters to estimate in the noisy-ability case, and fourteen parameters in noisy-selection case.<sup>37</sup>

To identify these structural parameters we use a core set of fourteen empirical moments, which we obtain from data on applications and test scores. We then show that our results are robust to the introduction of two additional moments based on self-reported beliefs data, which provide additional information on jobseeker uncertainty. For the core set of moments, we first use the control application rate and the control average and standard deviation of applicant ability (three moments). Further, we use the treatment effects on application rates and on average applicant ability of the two

---

<sup>37</sup>For noisy-ability case, these parameters are:  $\mu_{Tl}, \mu_{Cl}, \sigma_{Tl}, \sigma_{Cl}, \sigma_{TCl}, \mu_{Th}, \mu_{Ch}, \sigma_{Th}, \sigma_{Ch}, \sigma_{TCh}, \tau, \tau^w$  and  $a$ . For the noisy-selection case, these parameters are:  $\mu_{Tl}, \mu_{Cl}, \sigma_{Tl}, \sigma_{Cl}, \sigma_{TCl}, \mu_{Th}, \mu_{Ch}, \sigma_{Th}, \sigma_{Ch}, \sigma_{TCh}, \tau, \tau^w, \mu_a$  and  $\sigma_a$ .

interventions (four moments). For the high-wage group, we use the demediated treatment effect on application rates, which nets out the indirect impacts of this intervention on applications through changes in beliefs and salience (see Section IV.E for a full discussion of demediation). We compute these moments separately for low and high- $B$  jobseekers, giving us a total of fourteen core moments. Jobseekers’ average belief about the probability of being offered the job, for the low and high- $B$  group, give us the two additional moments.<sup>38</sup> We present analytical formulas for all moments in Appendix D.

The parameters are jointly identified by these empirical moments. The basic intuition for identification is as follows. The six control group moments — the average and standard deviation of applicant ability and the application rate — describe the truncated distribution of  $T$  among applicants. These moments enable us to identify the mean of application costs ( $\mu_C$ ) — a key driver of the application decisions — and, conditional on this, the mean and standard deviations of ability ( $\mu_T, \sigma_T$ ) among all jobseekers. Then, the treatment effects on application rates identify the severity of the shocks  $\tau$  and  $\tau^w$  and the standard deviations of costs  $\sigma_{Ch}$  and  $\sigma_{Cl}$ . Further, the treatment effects on applicant ability identify the covariance between cost and ability, as shown in Proposition 2 and Proposition 3. Finally, perceived selectivity is identified by the treatment effects on ability and by jobseekers’ forecasts of the probability of being offered the job. Tables A45 and A46 summarise these intuitions.

We proxy ability  $T$  with the score on the Raven test and then study whether our findings change if we use the cognitive ability score instead. Further, to measure  $B$  — the value of the experiment’s job net of outside options — we first obtain a post-LASSO forecast of the wage a jobseeker can expect to earn given their observables. We then calibrate  $B$  using the wage paid by the experiment’s job in the control condition, and an informed assumption about the probability of finding an alternative job at the jobseeker’s individual market wage. We discuss this calibration in detail in Appendix F. For the estimation, we discard individuals with a negative  $B$  since our model is only specified for the case where  $B > 0$ . We split the remaining individuals (about 61 percent of the sample) at the median level of  $B$ . On average, an individual in the high- $B$  group gets a net, discounted benefit from the experiment’s job of about 594 ETB (27.3 USD). For the low- $B$  group, the benefit is about 407 ETB (18.7 USD). A possible drawback of this strategy is that the residual variation in the value of the job that is not captured by the empirical measure of  $B$  may inflate our estimates of application costs and of the cost-

---

<sup>38</sup>We elicit beliefs about the probability of being offered the experiment’s job retrospectively. During the second phone call, we ask subjects to report the belief that they held at the time of deciding whether to apply to the position or not (that is, shortly after the first phone call). We show that these beliefs are not affected by the interventions in Table A39 in the Online Appendix.

ability correlation. We explore this point empirically in Section [V.D](#), by studying the robustness of our key results to estimation strategies that leverage additional variation in  $B$ , for example, by grouping individuals by the tercile in the distribution of  $B$  that they belong to.<sup>39</sup>

To estimate the model we use a classical minimum distance estimator ([Wooldridge, 2010](#)). We save the empirical moments in a vector  $\mathbf{m}$ . For a parameter vector  $\boldsymbol{\theta}$ , we solve the model and calculate the simulated moments  $\mathbf{m}_S(\boldsymbol{\theta})$ . We update  $\boldsymbol{\theta}$  in order to solve:

$$(9) \quad \hat{\boldsymbol{\theta}} = \min_{\boldsymbol{\theta}} [\mathbf{m}_S(\boldsymbol{\theta}) - \mathbf{m}]' \cdot J(\mathbf{m})^{-1} \cdot [\mathbf{m}_S(\boldsymbol{\theta}) - \mathbf{m}],$$

$J(\mathbf{m})$  is a diagonal matrix that contains the variance of each moment, ensuring that more precisely estimated moments get a greater weight in estimation. In line with the recent literature, we use this simple weighting matrix instead of the theoretically optimal weighting matrix, which may suffer from small sample bias ([Altonji and Segal, 1996](#)). We calculate  $J(\mathbf{m})$  using a bootstrap with 1,000 replications. We perform inference using a bootstrap (keeping  $J(\mathbf{m})$  fixed). We include the estimation of  $B$  and the demediation procedure in all bootstrap exercises.

Table [4](#) presents our main structural results. We include two sets of estimates for the noisy-ability case, obtained with the core set of moments and with the extended set of moments. We also include two sets of estimates for the noisy-selection case, again obtained with the two different sets of moments. In what follows, we first describe the empirical fit of these four versions of the model, and then discuss the parameter estimates and their robustness.

< [Table 4 here.](#) >

## V.B Model fit

We generally obtain a good fit between empirical and simulated moments. We show this in [Figure 7](#) below (and report the underlying moments in [Tables A48-A51](#) in the online Appendix). In both the noisy-ability and the noisy-selection cases, the simulated

---

<sup>39</sup>These strategies capture variation in  $B$  that is generated by market wages. On the other hand, we are unable to capture variation in  $B$  that comes from heterogeneity in the value of non-work time, as we do not have individual-level data on the likely determinants of this variable (e.g. childcare, access to government welfare programs, etc.). An additional point to note is that our interventions may have an impact on jobseekers' beliefs about market wages. In [Table A38](#), we presented evidence suggesting this may be the case for the wage intervention, but not for the application incentive intervention. This means that the shock  $\tau^w$  that we estimate captures the net effect of two countervailing forces: a direct impact on the value of the experiment's job (which raises the value of  $B$ ) and an indirect impact on the value of jobseekers' perceived outside options (which lowers the value of  $B$ ).

moments fit the empirical moments tightly. Simulated application rates are generally within one percentage point of the true empirical moment. The mean and standard deviation of the Raven test in the control group are matched almost exactly. Finally, we also fit fairly precisely the treatment effect on applicant ability. Overall, the noisy-ability and noise-selection cases generate similar goodness-of-fit statistics, which we report in the last row of Table 4.

< **Figure 7 here.** >

Further, the two versions of the model that do not target beliefs are able to reproduce jobseekers' substantial overconfidence about the probability of being offered the job. Among individuals with  $B > 0$ , the average forecast of this probability is 47 percent. This is outside a reasonable range, since the employer hires one person approximately every 115 applicants. Importantly, when we do not explicitly target these beliefs, our estimates imply a similar high degree of overconfidence. When ability is known, the implied average forecast of the probability of being offered the job is 64 percent. When ability is noisy, the average forecast is 35 percent. In addition, when we explicitly target these beliefs, both models match them very closely.

The estimated model also replicates a key non-targeted pattern in the data: in the control group, applicants have on average lower ability than non-applicants (as shown in Figure A4 for GPA, the only measure of ability that we observe for all jobseekers). This pattern intuitively suggests that the marginal applicant is better than the average applicant and hence, that the employer can attract better applicants by intervening on either the cost or the benefit margin. The estimated model reproduces this key pattern for both the noisy-ability and the noisy-selection case.

Finally, the elasticity of the simulated moments with respect to the model parameters broadly supports the intuition for identification presented above. As in [Kaboski and Townsend \(2011\)](#) and [Lagakos et al. \(2017\)](#), we first compute all moments using the estimated parameters. We then shock by one percent the value of each parameter at a time, and compute the percent change in the simulated moments. This illustrates the drivers of identification close to the minimum of the estimation problem. We report the results in Tables A52-A55 in the Appendix. The estimated elasticities are generally consistent with the intuitions for identification discussed above. For example, Table A52 shows that the elasticity of the treatment effect of application incentives on applicant ability with respect to the cost-ability covariance is 3.8 (a 1 percent change in the covariance leads to a 3.8 percent change in the simulated moment). For the moments that describe application rates or ability in the control group, the elasticity with respect

to this parameter is much lower. Tables A53 and A55 also illustrate how jobseekers' beliefs help identify perceived selectivity: these moments are most responsive to changes in the selectivity parameters.

## V.C Parameter estimates

### V.C.1 Noisy-ability case

Our estimates for the noisy-ability case indicate that application costs are large, heterogeneous, and positively correlated with ability. We report these estimates in the first two columns of Table 4. Column (1) reports the results obtained using the core set of empirical moments. For the high-value group, the mean of application costs is 217 ETB. This amounts to 13.5 percent of the monthly salary offered to individuals in the control group, or to about 36.5 percent of the estimated value of the job. For the low-value group, mean costs are about 136 ETB, or 8.5 percent of the salary and 33.5 percent of the value of the job. We also estimate that application costs have a large dispersion, in both groups. The standard deviation of application costs is about 245 ETB for the high  $B$  group and 193 for the low  $B$  group.

The correlation between ability and costs is 0.57 for the high  $B$  group and 0.64 for the low  $B$  group. These estimates imply a large increase in ability as we move along the cost distribution. For example, a high- $B$  jobseeker with costs one standard deviation above the mean has a Raven score that is about 6.4 scores higher than the average jobseeker (a 15 percent increase). Using the average Mincerian return to cognitive ability reported in Section I, we estimate that the value of this additional ability would be 146 ETB per month.

When we introduce jobseekers' beliefs, we obtain similar estimates of the cost-ability correlation and somewhat higher estimates of average costs (column (2) of Table 4). The estimated selectivity threshold is lower, which helps fit jobseekers' beliefs. All other parameters estimates are qualitatively unchanged.

### V.C.2 Noisy-selection case

Our estimates for the noisy-selection case confirm that application costs and the cost-ability correlation are substantial. We report these estimates in the last two columns of Table 4. We estimate correlation coefficients that are consistent with those of the noisy-ability case. However, mean application costs are substantially larger. When we use the core set of moments, we estimate costs of 273 ETB and 383 ETB. When we bring in the additional moments, we estimate costs of 206 and 286 ETB.

Uncertainty about selectivity is very large. As discussed in Section II, this uncertainty moderates the size of the “standard intuition” effect. We can show this by computing the increase in average applicant ability that we would observe if the correlation between cost and ability was set to zero (and all other parameters were equal to those reported in column (3) of Table 4). Thus, in this simulation, the treatment effect on ability is driven entirely by the “standard intuition” effect. We calculate that, for the high- $B$  group, the drop in average ability would correspond to 0.002 correct answers on the Raven test. This is less than 1 percent of the effect on ability that we obtain if we use the estimated positive value of the correlation, which suggests that the screening role of application costs is limited in our context.

### V.C.3 Plausibility checks

To explore the plausibility of our application cost estimates, we leverage jobseekers’ own reports about the monetary and time costs of other job applications.<sup>40</sup> The average monetary expenditure for one application is 50 ETB.<sup>41</sup> Further, jobseekers report that the average time required for one application is about 5 hours — which is broadly in line with the time required to apply for the experiment’s position. We estimate that the value of this time is between 56 ETB and 89 ETB.<sup>42</sup> The sum of monetary and time costs is thus between 106 and 139 ETB. This should be seen as a lower bound to the total cost of the application since it does not include non-material costs (related to stress, attention, effort, etc.). Given this, our structural estimates for the noisy-ability case (136

---

<sup>40</sup>This data was collected during the second phone interview. The question we asked referred to applications to positions other than the experiment’s job, in the 30 days between the two phone interviews. While most jobseekers make at most one job application in this period, a small group sends a large a number of applications. This last group also spends much less money and time per application, which complicates the interpretation of these cost values. Thus, in the analysis that follows, we use the average figures reported by jobseekers that make only one application.

<sup>41</sup>This figure is likely to reflect the multiple trips jobseekers make to complete the application process (in-person application, written test, interview) as well as printing CVs and other one-off expenses. It is thus not an estimate of the cost of a single day of commuting.

<sup>42</sup>The opportunity cost of time is high for people in our sample, who may not have a formal, stable job, but often have access to casual, informal income-generating opportunities. These opportunities are typically as remunerative, on an hourly wage basis, as entry-level formal jobs in the bottom part of the wage distribution. Thus, if we value the time spent on the application at the rate of the salary offered in the low wage condition of our experiment (assuming a 5-days, 40-hours working week and 2 statutory days off per month), the opportunity cost of time is about 56 ETB. This goes up to 89 ETB if we assume that a whole day worth of income-generating opportunities is lost — a reasonable assumption since one may have to forgo, say, a day of casual work at a construction site in order to spend 5 hours at the application centre.

and 217 ETB, with the core moments) are broadly consistent with these figures. On the other hand, the estimates for the noisy-selection case are noticeably larger.

Second, to assess the plausibility of our correlation estimates, we run a simple simulation based on the selection patterns we documented using the high-frequency panel. This panel tracks one cohort of jobseekers over the course of one year. For our simulation, we assume that the labour market is composed of multiple cohorts that live for three years, shrinking in size by 20 percent per year. Further, we assume that the cost-ability correlation in each cohort evolves as in the cohort that we observe in the panel data. Given this set-up, we find an implied steady-state correlation between cost and ability of about .35, which confirms that a substantial correlation is plausible in this labour market. However, this benchmark is lower than our structural estimates of the correlation parameter (which are between 0.57 and 0.64). This may be due to the fact that our proxies of application costs are noisy, which would bias the simulation downwards. Alternatively, our structural estimates may be inflated by variation in the value of the job that is not captured by our empirical measure of  $B$ . We explore this point in the robustness section below.

## V.D Robustness

The results on application costs and their correlation with ability are robust to the use of several alternative estimation strategies. For this robustness analysis, we focus on the noisy-ability case, which gives the most conservative estimates. We perform the following checks: (i) we use the cognitive ability score as opposed to the Raven score to measure ability, (ii) we let, in turn,  $a$ ,  $\tau^w$ , and  $\tau$  differ by  $B$  group, (iii) we predict  $B$  using an OLS model instead of the post-LASSO estimator and (iv) we allow for three types of  $B$  in the population. We report all key results in Table A56.

First, we find that, when we use the cognitive ability score to measure ability, we estimate very similar levels of costs and a higher cost-ability correlation. These results, reported in the first column of Table A56, suggest that our structural findings are not driven by our particular measure of ability. Second, we find that if we relax the assumption that  $a$ ,  $\tau$ , and  $\tau^w$  do not vary across groups, we estimate values for the high and low- $B$  groups that are qualitatively very similar to each other, as shown in columns (2)-(4). Third, we study two models designed to reduce the residual variation in  $B$ , either by splitting the distribution of  $B$  more finely, or by forecasting market wages using an OLS estimator (which extracts information from a larger set of covariates compared to the Post-LASSO estimator). When we split  $B$  into three groups, application costs range from 154 ETB to 270 ETB and the correlation is between 0.48 and 0.76. Further, when

we use the OLS estimator, the lowest estimate of application costs is 115 ETB and the lowest estimate of the cost-ability correlation is 0.51. Thus, under both strategies, we estimate large application costs and a strong cost-ability correlation. While this exercise cannot fully rule out that our estimates may be inflated by residual variation in the value of the job, it provides evidence of the qualitative robustness of our findings when this residual variation is plausibly reduced.

## **V.E The returns of the interventions and policy simulations**

In this last subsection, we assess the returns of the interventions and of two counterfactual policies. Both interventions attract applicants with higher cognitive ability and thus enable the employer to make better hires. Since cognitive ability is a strong predictor of productivity and the wage is fixed to the level that was originally posted, this generates an expected stream of benefits for the employer. Each intervention also entails two types of costs. First, the employer has to pay the direct cost of the intervention (the cost of the incentive or the higher wage). Second, the employer has to employ staff time to review the additional applications.

We calibrate our cost and benefit calculations in order to assess the effect of the interventions on a single, typical employer recruiting a clerical worker in this market. For these calibrations, we use the data that we collected from local firm managers. First, we quantify recruitment costs using managers' assessment of the time required to review one more application. On average managers report that this requires about one hour of work.<sup>43</sup> We price this hour at the median salary of the HR staff who review applications in these firms. Second, we calibrate the number of applicants in the control group and the number of jobs on offer using the average of these variables among the firms in our sample. Third, we compute worker turnover rates and use these to assess the expected number of months that the worker will spend in the firm (expected tenure is 45 months; we assume, conservatively, that the high wage is paid only for three months). Finally, we calibrate the productivity gains from higher worker ability using the Mincerian regressions on local labour market data that we presented in Section I. We summarise these assumptions and give additional details on our calculations in Appendix D.

We design two counterfactual policies that reduce the upfront costs of the application incentive. One drawback of this intervention is that the employer subsidises a large group of infra-marginal individuals who would have applied for the job in the absence

---

<sup>43</sup>We also ask whether there are any financial costs involved in reviewing one more application. The great majority of managers report that this is not the case. The majority of financial costs are fixed costs, such as advertisement costs.

of the incentive. Further, while most of the increase in ability is driven by the high- $B$  group, the incentive is offered to both groups. To decrease transfers to infra-marginal and to low-ability applicants we propose two alternative policies: (i) an application incentive that is targeted on the basis of observable demographics (we experiment, in turn, with offering the incentive only to women and only to individuals below the median age),<sup>44</sup> (ii) an application incentive that is offered only to the applicants who score above a threshold in the test (we set this threshold to the level that fills the positions on offer in expectation, so in practice under this policy the incentive is offered to all hires). These interventions reduce transfers to infra-marginal individuals by exploiting, in turn, the observable information available to employers and the information available to workers. One important caveat with respect to the second counterfactual intervention is that our estimates do not take risk aversion into account. Under this intervention, jobseekers are uncertain about whether they will qualify for the incentive. If risk aversion is large, the incentive will be less attractive than in our simulations. Further, if it is negatively correlated with ability, the positive effect on the quality of the pool of applicants will be less pronounced than what we predict. The returns that we estimate for this intervention should thus be seen as an upper bound of the true returns.<sup>45</sup>

< **Table 5 here.** >

We present these results in Table 5 below. First, we find that the application incentive has a positive internal rate of return (IRR) of about 9.8 percent. This is above market interest rates (which were about 5 percent in Ethiopia at the time of the experiment), and in line with the hurdles rates commonly reported by firms.<sup>46</sup> However, the IRR

---

<sup>44</sup>To estimate these counterfactuals, we estimate the structural model two more times, using new empirical moments obtained by splitting the sample by age and by gender. We report these new parameter estimates in Table A57 and A58, and the related moments in Table A59 and A60. The model that explores heterogeneity by age fits the data well. On the other hand, the model that explores heterogeneity by gender has a poorer fit. This is driven by the fact that, for the incentive intervention, the model predicts an increase in the ability of female applicants that is only about one fifth of the impact observed in the data. In our IRR calculations, we thus conservatively augment the simulated impact of this intervention, reducing the gap between the empirical and simulated moment by 25 percent. Our estimates are hence likely to be a lower bound of the true returns of this counterfactual intervention.

<sup>45</sup>For both counterfactual policies, we also assume that treatment does not change participants' beliefs about the probability of getting the job. This is justified by the fact that, in the experiment, the incentive intervention does not have an effect on these beliefs. We are however unable to test whether the targeted policies that we examine in the counterfactual analysis may impact beliefs differently compared to the untargeted incentive that we have implemented in the field.

<sup>46</sup>We are not aware of data on the hurdle rates used by firms in Ethiopia or in countries with similar

estimate is somewhat imprecise and the bootstrapped confidence interval has a negative lower bound. Thus, while the intervention is attractive in expectation, we cannot reject that the IRR is different from zero (or different from a meaningful negative level). Second, we find that all counterfactual incentive schemes have a large impact on the return of the intervention. Targeting the incentive to women or to jobseekers below the median age raises the IRR to 86 percent and 47 percent, respectively. Both of these schemes substantially reduce the cost of the intervention, as they restrict the number of applicants that are eligible for the incentive. At the same time, they generate large gains in applicant ability, since they target the incentive on the demographics that drive the impacts of the intervention. Thus, they raise the IRR considerably; and the bootstrap exercise suggests that these counterfactual estimates are significantly different from zero. Third, when the incentive is offered to all hires, the cost of the intervention decreases dramatically, and the IRR goes up to 382 percent. However, as discussed above, this large IRR should be interpreted as an upper bound, as it does not consider the role of risk aversion.

Our IRR estimates do not measure whether the intervention generates overall welfare gains. The fact that high-*B* individuals drive the impacts on ability (Table A27) suggests that application incentives may have positive effects on the allocation of talent, favouring those high-ability jobseekers who stand to gain the most from being offered the job. However, we do not have access to the comprehensive labour market data that would be required to credibly quantify these allocative gains. In particular, we do not have information on how other firms are affected by the policy or on workers' long-term outcomes. We thus do not provide an estimate of welfare effects.

## VI Discussion

In this section we discuss two important questions that follow from our findings. First, we study whether employers value cognitive ability and whether they are aware of the impacts of application incentives on the quality of the applicant pool. Second, in order to shed light on the external validity of our results, we present evidence from other contexts on application costs and their correlation with applicant ability.

---

macroeconomic conditions. A recent survey by the Bank of England finds that most firms in the UK adopt hurdle rates between 5 and 15 percent (Saleheen et al., 2017).

## VI.A Do managers value cognitive ability and are they aware of the benefits of application incentives?

We address these questions through a second novel experiment that studies the preferences and expectations of managers at firms recruiting for clerical positions (see section I for a description of how these firms are sampled). This experiment enables us to explore two possible reasons why firms in Addis Ababa do not use application incentives. First, firm managers may not value general cognitive ability and thus may not rank applicants from the incentive group above control applicants. Alternatively, managers would like to recruit workers with higher cognitive ability, but do not expect that the application incentive will attract these workers.

In the first task, we study whether firm managers rank treatment group applicants above control applicants. To incentivise this task, we offer to invite one person from our sample of applicants to make a new job application at the manager's firm. The manager can determine who this person will be by ranking the standardised CVs of three selected applicants.<sup>47</sup> We sample one applicant from each experimental group. At this point of the experiment, however, the manager has not been informed about the two interventions nor about how the three applicants have been selected. On the CVs, we report applicants' education, age, work experience, GPA and the results from the Raven and conscientiousness tests (Figure A15 shows a sample CV). We select triplets of applicants that reproduce as closely as possible the average differences in these characteristics between groups.<sup>48</sup> After the manager ranks the CVs, we randomly draw two of the three CVs and invite the person with the higher rank to make an application at the manager's firm. This last feature ensures that the manager is incentivised to report truthfully her preferences over the three candidates.

We find that both interventions improve the quality of the applicant pool in the eyes of local firm managers. We show this result in Table 6 using a series of linear probability

---

<sup>47</sup>Kessler et al. (2019) propose a similar methodology — called Incentivised Resume Rating — to elicit employer preferences. There are two main difference between their methodology and ours: (i) they use fictitious CVs while we rely on real CVs, and (ii) in their design, employers assess CVs by reporting a cardinal score, while in our design employers report a rank.

<sup>48</sup>We do this in three steps. First, we generate a list of all possible triplets. Second, for each triplet, we calculate the distance between (i) the differences in ability among the applicants in the triplet and (ii) the average differences in ability among their respective experimental groups (with a simple sum-of-square-differences statistic and three measures of ability: GPA, Raven score and conscientiousness). Third, we choose the sixteen triplets that have the minimum value of this distance statistic. We then randomly allocate a triplet to each manager. Across triplets, we also randomly change the order with which the candidates from the three experimental groups are presented.

models. In the first two columns, the dependent variable is a dummy for individuals who are ranked first. In the third column, the dependent variable is a dummy for being ranked first or second. We find that applicants from the incentive group are a significant 36.9 percent more likely to be ranked first than control applicants, and a significant 37.4 percentage points more likely to be ranked first or second. In column two, we only consider applicants from the control and incentive groups. We find that incentive group applicants are ranked above control group applicants about 70 percent of the times.

< **Table 6 here.** >

In the second task, we test whether managers can predict the effects of the application incentive. We first give managers detailed information on the experiment and then ask them to forecast the impacts of the application incentive on application rates and applicant ability (as measured by the Raven test). To measure ability at different points of the distribution, we obtain forecasts of (i) the average Raven score and (ii) the average Raven score among the 100 highest-scoring applicants. Further, before forecasts are made, we disclose the application rates and Raven test scores of applicants in the control and high wage groups in order to anchor managers' priors on the correct level of these variables. We reward managers for the accuracy of one randomly drawn forecast.

We find that managers make considerable forecasting errors and generally underestimate the impacts of the application incentive on applicant ability. In Figure 8 we plot the distribution of managers' forecasts. On average, managers expect that the application incentive will increase application rates and decrease applicant ability. In particular, they predict that performance on the Raven test will fall by about one correct answer, both at the mean and at the top of the distribution. In reality, performance on the Raven test improves by about one correct answer in both cases. Overall, about 66 percent of managers underestimate the level of cognitive ability of the applicants in the incentive group.

In sum, the evidence in this section shows that managers value applicant cognitive ability, but are unaware of the cognitive ability gains generated by application incentives.

< **Figure 8 here.** >

## VI.B External validity

We find several pieces of evidence suggesting that job search and applications costs are large in a number of contexts. In China, [Chang \(2009\)](#) reports qualitative evidence that manufacturing workers do not apply for attractive jobs when the application centre is too far from their place of residence. In Ghana, [Hardy and McCasland \(2017\)](#) provide evidence that entry fees for apprentices create labor constraints for small firms. In South Africa, [Abel et al. \(2017\)](#) show that jobseekers apply to fewer jobs per week than they would like to. A psychological intervention is only able to reduce this intention-behaviour gap by a small amount, suggesting that other constraints such as credit constraints may be responsible for jobseekers' failure to meet their job application targets. In OECD countries, employers often reimburse applicants for travel expenses incurred for job interviews. In Germany, this is mandated by the Civil Code.<sup>49</sup> In other markets, this emerges as a stable equilibrium outcome. An example of this is the economics job market, where most universities reimburse applicants for 'flyout' expenses. In all of these contexts, job application costs emerge as a salient constraint for jobseekers.

We also find observational evidence suggesting that these costs may be positively correlated with worker ability in several settings. In India, [Choudhury and Khanna \(2014\)](#) leverage micro-data from a large technology firm to show that hires from remote small towns (who face larger search costs for formal work) have higher ability and are more productive than hires from large cities (who face lower search costs). In the US, [Paserman \(2008\)](#) structurally estimates that, among unmarried jobseekers with low-wage work experience, high-ability jobseekers face search costs that are almost 5 times higher than those of low ability jobseekers.<sup>50</sup> Finally, there is evidence of similar selection dynamics in the context of college admissions. For example, [Black et al. \(2020\)](#) study the impacts of a policy granting public university access to all Texas students in the top 10 percent of their high-school class. They find that those top students from schools with poor university placement records who attend college thanks to the policy — a group that presumably faces large costs to access higher education — have better graduation rates than the students from the traditional "feeder" schools who are displaced by the policy. Overall, this evidence suggests that, across a variety of contexts, high-ability marginal applicants tend to face high application costs. In these settings,

---

<sup>49</sup>See §670 of the German Civil Code (BGB) and <https://www.euraxess.de/germany/information-assistance/work-permit/faq-working>.

<sup>50</sup>US jobseekers with low-wage employment experience are likely to be the most comparable group to the jobseekers in our sample. For jobseekers with experience in better-paid positions, [Paserman \(2008\)](#) finds a negative cost-ability correlation.

reducing application costs can improve both the efficiency and equity of the selection process.

One final important consideration is that the cost-ability correlation may be affected by different types of labour market shocks. For example, the introduction of low-cost online job portals may equalise individuals' ability to search for work, and thus eliminate the differential dynamic selection pattern we have documented. Furthermore, temporary negative shocks in demand may dampen the job-search advantage of low-cost individuals. These observations suggest that the mechanism we have uncovered may evolve over time.

## VII Conclusion

In a worker recruitment experiment in Addis Ababa, Ethiopia we show that employers can use application incentives to attract applicants with higher cognitive ability. We estimate a structural model of applications decisions and find that the positive effect of application incentives follows from the fact that application costs are large, heterogeneous and, surprisingly, positively correlated with jobseeker ability. Finally, using a high-frequency panel dataset on job-search decisions, we present evidence suggesting that the correlation between application costs and jobseeker ability is likely to be driven by a dynamic selection mechanism: low-cost, high-ability jobseekers find work faster than high-cost jobseekers with similar ability and thus quit job search earlier.

Our experimental evidence on how application costs affect firms' ability to recruit talented workers generates a number of leads for future research. A first natural question is how incentives interact with interventions that improve the quality of screening (Autor and Scarborough, 2008). Second, our study does not investigate the potential dynamic gains from relaxing labour constraints. If personnel ability is complementary to capital and technology (Bender et al., 2016), these gains could be large. More broadly, our results suggest that unequal access to labour markets can distort the selection and allocation of talent. This central insight can help design future labour market policy.

## References

- Abebe, Girum, Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, and Simon Quinn**, “Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City,” *Unpublished*, 2020.
- Abel, Martin, Rulof Burger, Eliana Carranza, and Patrizio Piraino**, “Bridging the Intention-Behavior Gap? The Effect of Plan-Making Prompts on Job Search and Employment,” 2017.
- Acharya, Avidit, Matthew Blackwell, and Maya Sen**, “Explaining Causal Findings Without Bias: Detecting and Assessing Direct Effects,” *American Political Science Review*, 2016, 110 (3), 512–529.
- Alatas, Vivi, Ririn Purnamasari, Matthew Wai-Poi, Abhijit Banerjee, Benjamin A Olken, and Rema Hanna**, “Self-Targeting: Evidence from a Field Experiment in Indonesia,” *Journal of Political Economy*, 2016, 124 (2), 371–427.
- Alonso, Ricardo**, “Recruiting and Selecting for Fit,” *Unpublished*, 2018.
- Altonji, Joseph G and Lewis M Segal**, “Small-Sample Bias in GMM Estimation of Covariance Structures,” *Journal of Business & Economic Statistics*, 1996, 14 (3), 353–366.
- Anderson, Michael L**, “Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects,” *Journal of the American statistical Association*, 2008, 103 (484).
- Ashraf, Nava, James Berry, and Jesse M Shapiro**, “Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia,” *American Economic Review*, 2010, 100 (5), 2383–2413.
- , **Oriana Bandiera, Scott S Lee et al.**, “Losing Prosociality in the Quest for Talent? Sorting, Selection, and Productivity in the Delivery of Public Services,” *Unpublished*, 2018.
- Attanasio, Orazio and Katja Kaufmann**, “Educational Choices, Subjective Expectations, and Credit Constraints,” *NBER Working Paper No. 15087*, 2009.
- Augenblick, Ned, Muriel Niederle, and Charles Sprenger**, “Working over Time: Dynamic Inconsistency in Real Effort Tasks,” *The Quarterly Journal of Economics*, 2015, pp. 1067–1115.

- Autor, David H and David Scarborough**, “Does Job Testing Harm Minority Workers? Evidence from Retail Establishments,” *The Quarterly Journal of Economics*, 2008, 123 (1), 219–277.
- **and Michael J Handel**, “Putting Tasks to the Test: Human Capital, Job Tasks, and Wages,” *Journal of Labor Economics*, 2013, 31 (S1), S59–S96.
- Balakrishnan, Uttara, Johannes Haushofer, and Pamela Jakiela**, “How Soon Is Now? Evidence of Present Bias from Convex Time Budget Experiments,” *Unpublished*, 2015.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul**, “Field Experiments with Firms,” *The Journal of Economic Perspectives*, 2011, 25 (3), 63–82.
- Banerjee, Abhijit V and Andrew F Newman**, “Occupational Choice and the Process of Development,” *Journal of political economy*, 1993, 101 (2), 274–298.
- **and Sandra Sequeira**, “Spatial Mismatches and Imperfect Information in the Job Search,” 2020.
- Barrett, Garry F and Stephen G Donald**, “Consistent Tests for Stochastic Dominance,” *Econometrica*, 2003, 71 (1), 71–104.
- Beaman, Lori, Niall Keleher, and Jeremy Magruder**, “Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi,” *Unpublished*, 2013.
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen**, “High-Dimensional Methods and Inference on Structural and Treatment Effects,” *The Journal of Economic Perspectives*, 2014, 28 (2), 29–50.
- Belot, Michelle, Phillip Kircher, and Paul Muller**, “How Wage Announcements Affect Job Search Behaviour - A Field Experimental Investigation,” *Unpublished*, 2017.
- Bender, Stefan, Nicholas Bloom, David Card, John Van Reenen, and Stefanie Wolter**, “Management Practices, Workforce Selection and Productivity,” *NBER Working Paper 22101*, 2016.
- Benjamini, Yoav, Abba M Krieger, and Daniel Yekutieli**, “Adaptive Linear Step-up Procedures that Control the False Discovery Rate,” *Biometrika*, 2006, 93 (3), 491–507.
- Black, Sandra E, Jeffrey T Denning, and Jesse Rothstein**, “Winners and Losers? The Effect of Gaining and Losing Access to Selective Colleges on Education and Labor Market Outcomes,” 2020.

- Bó, Ernesto Dal, Frederico Finan, and Martín A Rossi**, “Strengthening state capabilities: The role of financial incentives in the call to public service,” *The Quarterly Journal of Economics*, 2013, 128 (3), 1169–1218.
- Botta, Fabiano, Valerio Santangelo, Antonino Raffone, Juan Lupiáñez, and Marta Olivetti Belardinelli**, “Exogenous and Endogenous Spatial Attention Effects on Visuospatial Working Memory,” *The Quarterly Journal of Experimental Psychology*, 2010, 63 (8), 1590–1602.
- Bruhn, Miriam and David McKenzie**, “In Pursuit of Balance: Randomization in Practice in Development Field Experiments,” *American Economic Journal: Applied Economics*, 2009, 1 (4), 200–232.
- Bryan, Gharad and Melanie Morten**, “Economic Development and the Spatial Allocation of Labor: Evidence from Indonesia,” *Unpublished*, 2015.
- Bureau of Labor Statistics**, “News Release Job Openings and Labor Turnover: June 2017,” 2017.
- Card, David, Jochen Kluve, and Andrea Weber**, “Active Labour Market Policy Evaluations: A Meta-Analysis,” *The economic journal*, 2010, 120 (548).
- Caselli, Francesco, Jacopo Ponticelli, and Federico Rossi**, “A New Dataset on Mincerian Returns,” *Technology differences over space and time*, 2014, pp. 108–118.
- Chamorro-Premuzic, Tomas and Adrian Furnham**, *The Psychology of Personnel Selection*, Cambridge University Press, 2010.
- Chang, Leslie T**, *Factory Girls: From Village to City in a Changing China*, Random House Digital, Inc., 2009.
- Choudhury, Prithwiraj and Tarun Khanna**, “Firms and Within Country Migration: Evidence from India,” 2014.
- Crawford, Vincent P, Miguel A Costa-Gomes, and Nagore Iriberry**, “Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications,” *Journal of Economic Literature*, 2013, 51 (1), 5–62.
- Crépon, Bruno and Gerard J Van den Berg**, “Active Labor Market Policies,” *Annual Review of Economics*, 2016, 8, 521–546.

- , **Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora**, “Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment,” *The Quarterly Journal of Economics*, 2013, 128 (2), 531–580.
- Deaton, Angus**, *The Analysis of Household Surveys: a Microeconometric Approach to Development Policy*, World Bank Publications, 1997.
- DellaVigna, Stefano**, “Psychology and Economics: Evidence from the Field,” *Journal of Economic Literature*, 2009, 47 (2), 315–372.
- , “Structural Behavioral Economics,” *NBER Working Paper no. 24797*, 2018.
- **and Devin Pope**, “What Motivates Effort? Evidence and Expert Forecasts,” *NBER Working Paper No. 22193*, 2016.
- **and Matthew Gentzkow**, “Uniform Pricing in US Retail Chains,” *NBER Working Paper No. 23996*, 2017.
- Deserranno, Erika**, “Financial Incentives as Signals: Experimental Evidence from the Recruitment of Village Promoters in Uganda,” *American Economic Journal: Applied Economics*, 2019, 11 (1), 277–317.
- Diamond, Adele**, “Executive Functions,” *Annual review of psychology*, 2013, 64, 135–168.
- Dohmen, Thomas J, Armin Falk, David Huffman, and Uwe Sunde**, “Interpreting Time Horizon Effects in Inter-Temporal Choice,” *Unpublished*, 2012.
- Duckworth, Angela L, Christopher Peterson, Michael D Matthews, and Dennis R Kelly**, “Grit: Perseverance and Passion for Long-Term Goals.,” *Journal of Personality and Social Psychology*, 2007, 92 (6), 1087.
- Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde**, “Global Evidence on Economic Preferences,” *Unpublished*, 2016.
- , – , – , **David Huffman, and Uwe Sunde**, “The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences,” *Unpublished*, 2016.
- Flory, Jeffrey A, Andreas Leibbrandt, and John A List**, “Do Competitive Workplaces Deter Female Workers? A Large-Scale Natural Field Experiment on Job Entry Decisions,” *The Review of Economic Studies*, 2014, 82 (1), 122–155.

- Franklin, Simon**, "Location, Search Costs and Youth Unemployment: Experimental Evidence from Transport Subsidies," *The Economic Journal*, 11 2017, 128 (614), 2353–2379.
- Gabaix, Xavier and David Laibson**, "Bounded Rationality and Directed Cognition," *Unpublished*, 2005.
- Galenianos, Manolis, Philipp Kircher, and Gábor Virág**, "Market Power and Efficiency in a Search Model," *International Economic Review*, 2011, 52 (1), 85–103.
- Gardner, Donald G and Jon L Pierce**, "The Core Self-Evaluation scale: Further Construct Validation Evidence," *Educational and Psychological Measurement*, 2010, 70 (2), 291–304.
- Gobillon, Laurent, Harris Selod, and Yves Zenou**, "The Mechanisms of Spatial Mismatch," *Urban studies*, 2007, 44 (12), 2401–2427.
- Hardy, Morgan and Jamie McCasland**, "Are Small Firms Labor Constrained? Experimental Evidence from Ghana," *Unpublished*, 2017.
- Hsieh, Chang-Tai and Enrico Moretti**, "Why do Cities Matter? Local Growth and Aggregate Growth," *NBER Working Paper No. 21154*, 2015.
- , **Erik Hurst, Charles I Jones, and Peter J Klenow**, "The Allocation of Talent and US Economic Growth," *NBER Working Paper No. 18693*, 2013.
- Imbert, Clement and John Papp**, "Short-term Migration Costs: Evidence from India's Employment Guarantee," *Unpublished*, 2016.
- John, Oliver P and Sanjay Srivastava**, "The Big Five Trait Taxonomy: History, Measurement, and Theoretical Perspectives," *Handbook of Personality: Theory and Research*, 1999, 2 (1999), 102–138.
- Kaboski, Joseph P and Robert M Townsend**, "A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative," *Econometrica*, 2011, 79 (5), 1357–1406.
- Kessler, Judd B, Corinne Low, and Colin D Sullivan**, "Incentivized Resume Rating: Eliciting Employer Preferences without Deception," *American Economic Review*, 2019, 109 (11), 3713–44.
- Koenker, Roger and Kevin Hallock**, "Quantile Regression: An Introduction," *Journal of Economic Perspectives*, 2001, 15 (4), 43–56.

- Kremer, Michael, Gautam Rao, and Frank Schilbach**, “Behavioral Development Economics,” *Unpublished*, 2018.
- Laajaj, Rachid and Karen Macours**, “Measuring Skills in Developing Countries,” *Unpublished*, 2017.
- Lagakos, David, Mushfiq Mobarak, and Michael E Waugh**, “The Welfare Effects of Encouraging Rural-Urban Migration,” *Unpublished*, 2017.
- Lefcourt, Herbert M**, *Locus of Control*, Academic Press, 1991.
- Machado, J. A. F., P.M.D.C Parente, and J. M. C. Santos Silva**, “QREG2: Stata Module to Perform Quantile Regression with Robust and Clustered Standard Errors,” 2011.
- Machado, Jose AF and JMC Santos Silva**, “Glejser’s Test Revisited,” *Journal of Econometrics*, 2000, 97 (1), 189–202.
- Malmendier, Ulrike and Geoffrey Tate**, “Behavioral CEOs: The Role of Managerial Overconfidence,” *The Journal of Economic Perspectives*, 2015, 29 (4), 37–60.
- Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao**, “Poverty Impedes Cognitive Function,” *Science*, 2013, 341 (6149), 976–980.
- Marimon, Ramon and Fabrizio Zilibotti**, “Unemployment vs. Mismatch of Talents: Reconsidering Unemployment Benefits,” *The Economic Journal*, 1999, 109 (455), 266–291.
- Mas, Alexandre and Amanda Pallais**, “Valuing Alternative Work Arrangements,” *NBER Working Paper No. 22708*, 2016.
- and —, “Labor Supply and the Value of Non-Work Time: Experimental Estimates from the Field,” *Unpublished*, 2017.
- McKenzie, David**, “How Effective are Active Labor Market Policies in Developing Countries? A Critical Review of Recent Evidence,” *The World Bank Research Observer*, 2017.
- , “Identifying and Spurring High-Growth Entrepreneurship: Experimental Evidence from a Business Plan Competition,” *American Economic Review*, 2017, 107 (8), 2278–2307.
- Mullainathan, Sendhil and Eldar Shafir**, *Scarcity: Why Having Too Little Means So Much*, Macmillan, 2013.

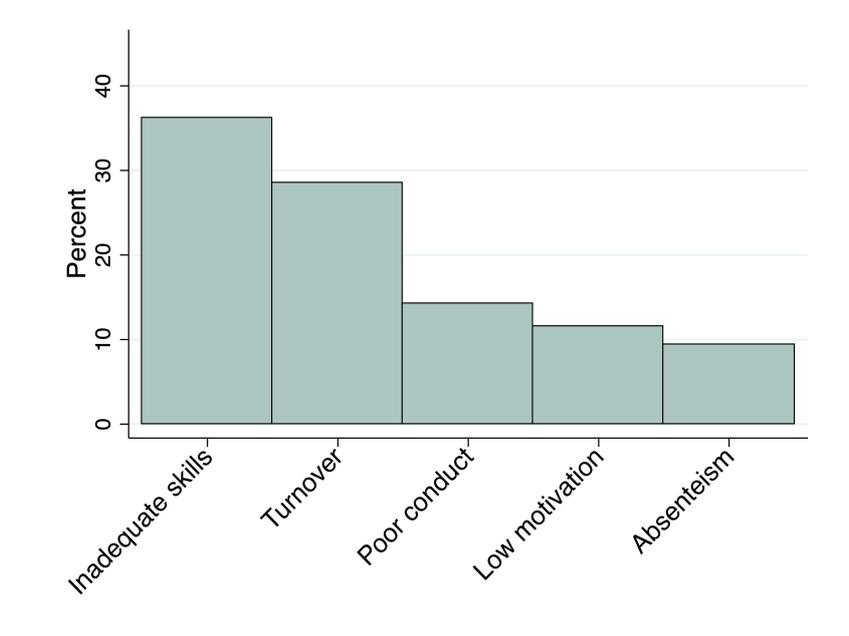
- Nagel, Rosemarie**, "Unraveling in Guessing Games: An Experimental Study," *The American Economic Review*, 1995, 85 (5), 1313–1326.
- Nichols, Albert L and Richard J Zeckhauser**, "Targeting Transfers Through Restrictions on Recipients," *The American Economic Review*, 1982, 72 (2), 372–377.
- Paserman, M Daniele**, "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation," *The Economic Journal*, 2008, 118 (531), 1418–1452.
- Phillips, David C**, "Getting to work: Experimental evidence on job search and transportation costs," *Labour Economics*, 2014, 29, 72–82.
- Porzio, Tommaso**, "Cross-Country Differences in the Optimal Allocation of Talent and Technology," 2017.
- Raven, John**, "The Raven's Progressive Matrices: Change and Stability over Culture and Time," *Cognitive psychology*, 2000, 41 (1), 1–48.
- Rogerson, Richard, Robert Shimer, and Randall Wright**, "Search-theoretic models of the labor market: A survey," *Journal of economic literature*, 2005, 43 (4), 959–988.
- Rosenberg, Morris**, *Conceiving the Self*, RE Krieger, 1986.
- Saleheen, Jumana, Iren Levina, Marko Melolinna, and Sr?an Tatomir**, "The Financial System and Productive Investment: New Survey Evidence," *Bank of England Quarterly Bulletin*, 2017, Q1.
- Santangelo, Valerio and Emiliano Macaluso**, "Visual Saliency Improves Spatial Working Memory via Enhanced Parieto-Temporal Functional Connectivity," *Journal of Neuroscience*, 2013, 33 (9), 4110–4117.
- Schmidt, Frank L and John E Hunter**, "The Validity and Utility of Selection Methods in Personnel Psychology: Practical and Theoretical Implications of 85 Years of Research Findings," *Psychological bulletin*, 1998, 124 (2), 262.
- Spinnewijn, Johannes**, "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs," *Journal of the European Economic Association*, 2015, 13 (1), 130–167.
- Vansteelandt, Stijn**, "Estimating Direct Effects in Cohort and Case–Control Studies," *Epidemiology*, 2009, 20 (6), 851–860.
- Weaver, Jeff**, "Jobs for Sale: Bribery and Misallocation in Hiring," *Unpublished*, 2016.

**Wooldridge, Jeffrey M**, *Econometric Analysis of Cross Section and Panel Data*, MIT press, 2010.

**Zenou, Yves**, *Urban Labor Economics*, Cambridge University Press, 2009.

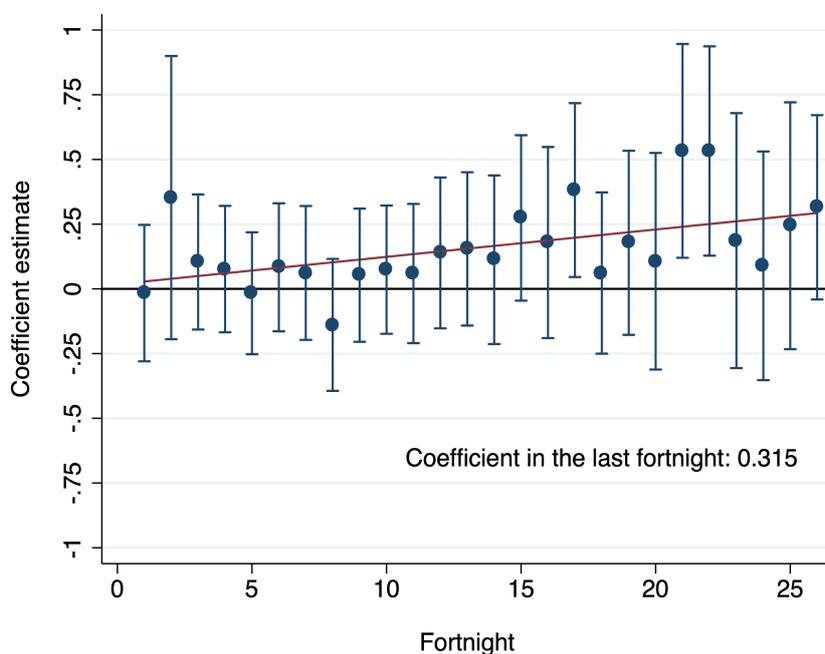
## Figures and tables for inclusion in the main text

Figure 1: Most important HR problem



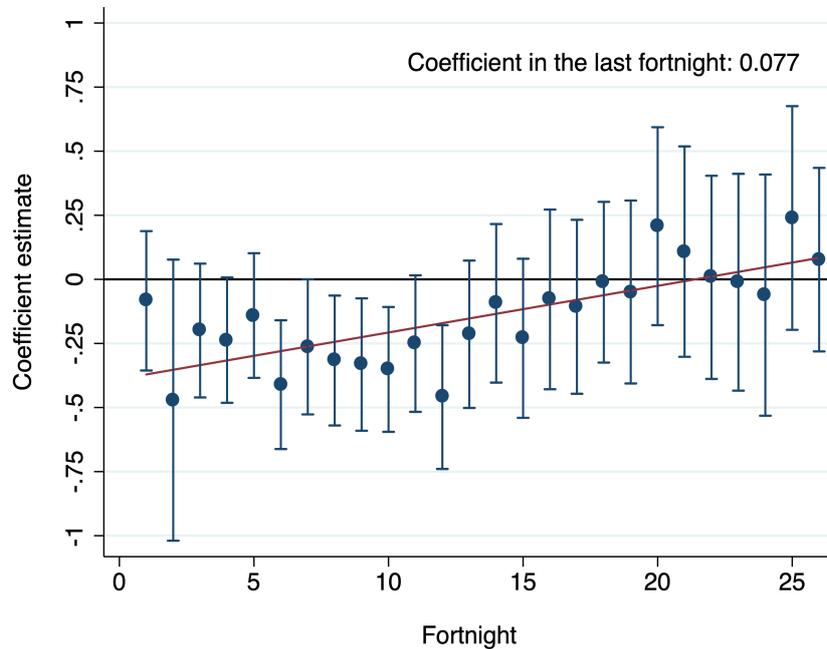
*Notes:* This figure reports data from our survey of firms hiring clerical workers. We report the distribution of managers' responses to the question 'What is the most important HR problem faced by your firm?'. Sample used: all managers.

Figure 2: Low savings and Raven test score among jobseekers by fortnight



*Notes:* In this figure, we plot the coefficients from a battery of regressions of standardised Raven test scores on a dummy for having below-median savings. We estimate a separate regression for each fortnight, each time restricting the sample to those individuals who are out of work and searching for employment in that particular fortnight. This regression thus estimates the correlation between savings and Raven test scores among the individuals that are actively looking for employment in a given fortnight. Changes in this correlation are driven by selection in and out of job search. The data comes from the high-frequency panel of [Abebe et al. \(2020\)](#). The jobseekers in this dataset are interviewed on the phone every two weeks for a period of one year. The same jobseekers also complete face-to-face baseline and endline interviews before and after the phone survey. For logistical reasons, the first three fortnights and the last fortnight of the phone survey have much fewer observations than the other fortnights and we thus drop them. We include face-to-face baseline and endline observations and consider them as the first and last fortnight of the panel. For each regression, we report the point estimate of the coefficient on the dummy variable for having below-median savings and a 90 percent confidence interval obtained using robust standard errors.

Figure 3: High distance and Raven test score among jobseekers by fortnight



*Notes:* In this figure, we plot the coefficients from a battery of regressions of standardised Raven test scores on a dummy for living at a distance from the city centre above the median in the sample. We estimate a separate regression for each fortnight, each time restricting the sample to those individuals who are out of work and searching for employment in that particular fortnight. This regression thus estimates the correlation between distance and Raven test scores among the individuals that are actively looking for employment in a given fortnight. Changes in this correlation are driven by selection in and out of job search. The data comes from the high-frequency panel of [Abebe et al. \(2020\)](#). The jobseekers in this dataset are interviewed on the phone every two weeks for a period of one year. The same jobseekers also complete face-to-face baseline and endline interviews before and after the phone survey. For logistical reasons, the first three fortnights and the last fortnight of the phone survey have much fewer observations than the other fortnights and we thus drop them. We include face-to-face baseline and endline observations and consider them as the first and last fortnight of the panel. For comparability with Figure 2, we restrict the sample to observations with non-missing savings. Results are qualitatively unchanged if we do not apply this restriction. For each regression, we report the point estimate of the coefficient on the dummy variable for living at above-median distance from the city centre and a 90 percent confidence interval obtained using robust standard errors.

Figure 4: Illustration of the effect of an application incentive  
 (noisy-selection case,  $\rho_z > 0$ )

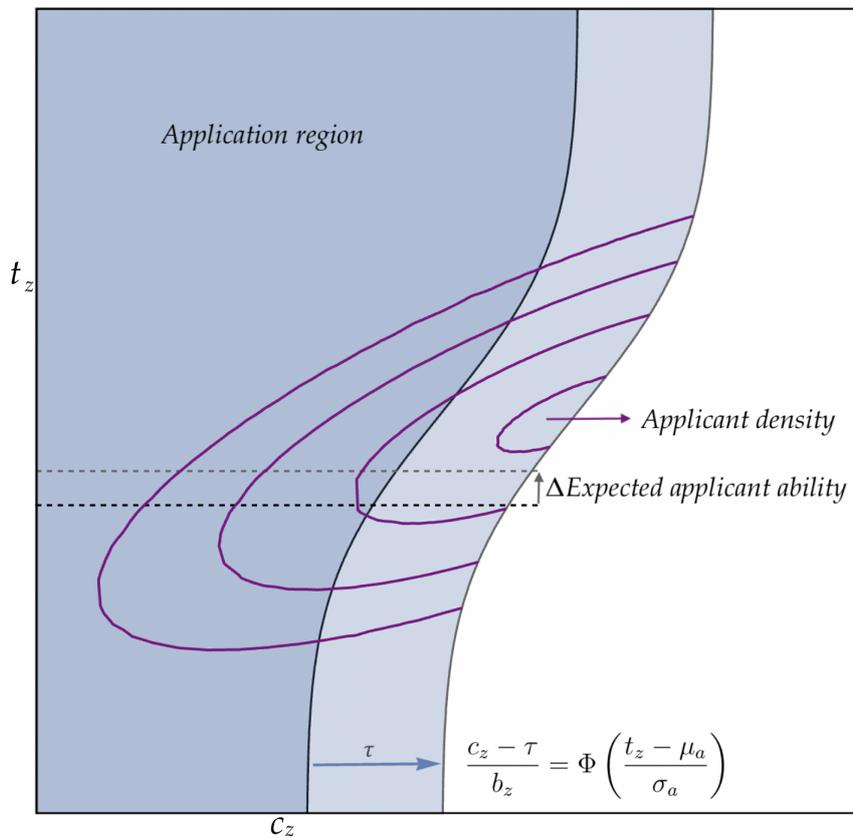


Figure 5: Illustration of the effect of an application incentive  
(noisy-ability case,  $\rho_z > 0$ )

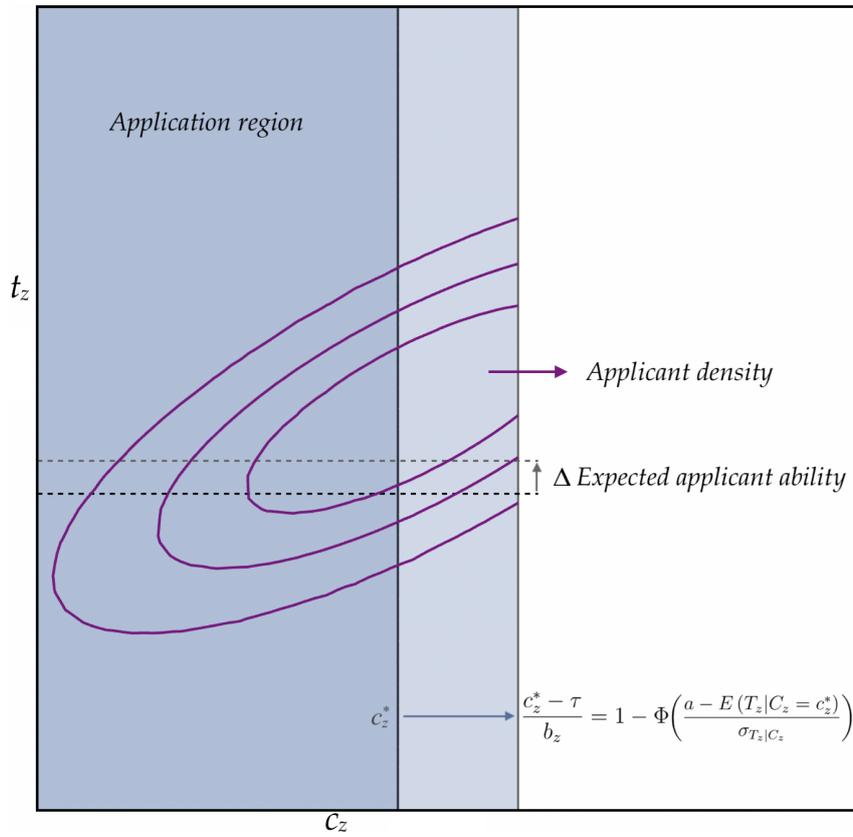
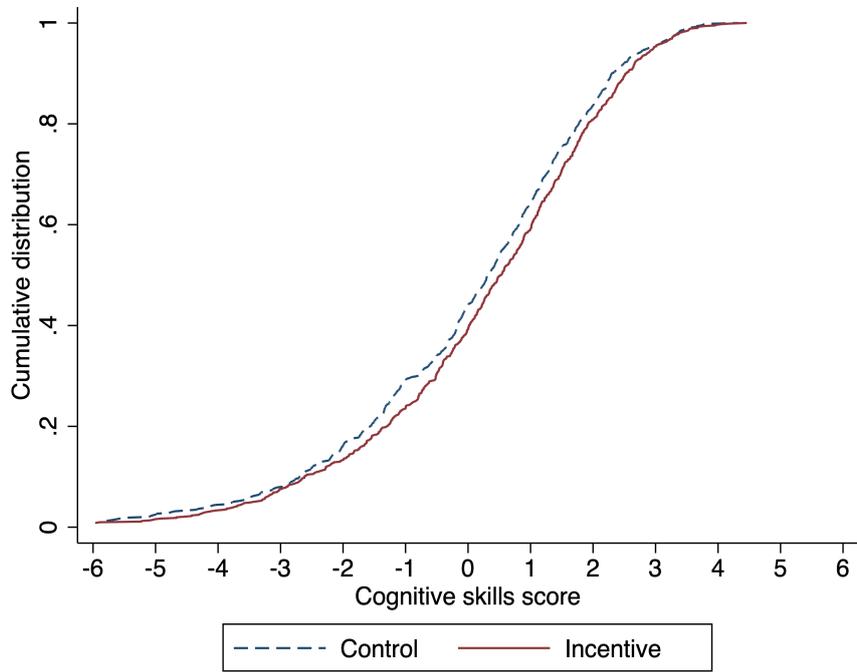
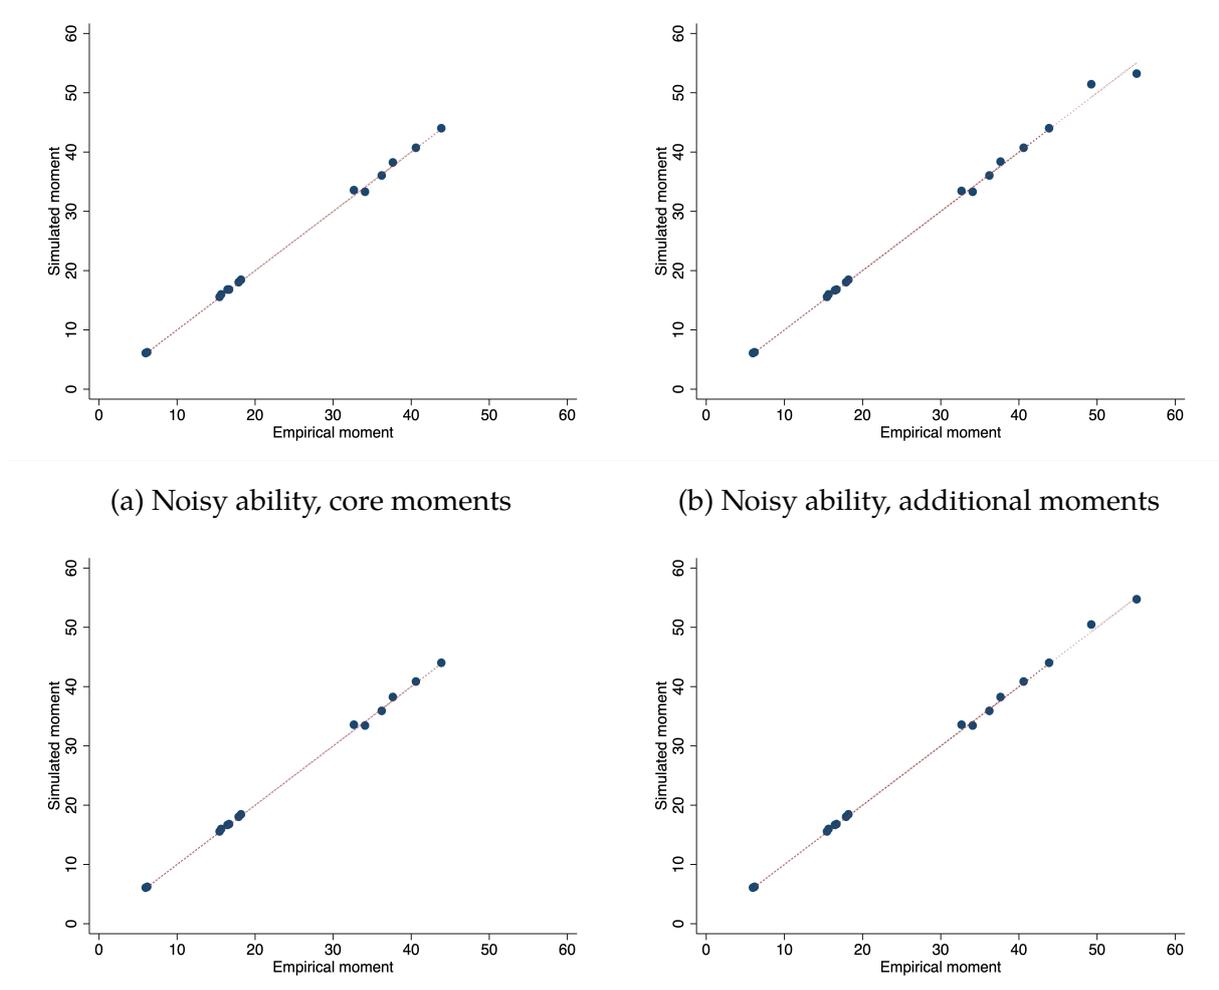


Figure 6: Impacts on the distribution of applicant cognitive ability  
Incentive treatment



*Notes:* This figure plots the distribution of cognitive ability among control and application incentive applicants in the experiment. Sample used: all applicants (incentive and control groups).

Figure 7: Moment fit



(a) Noisy ability, core moments

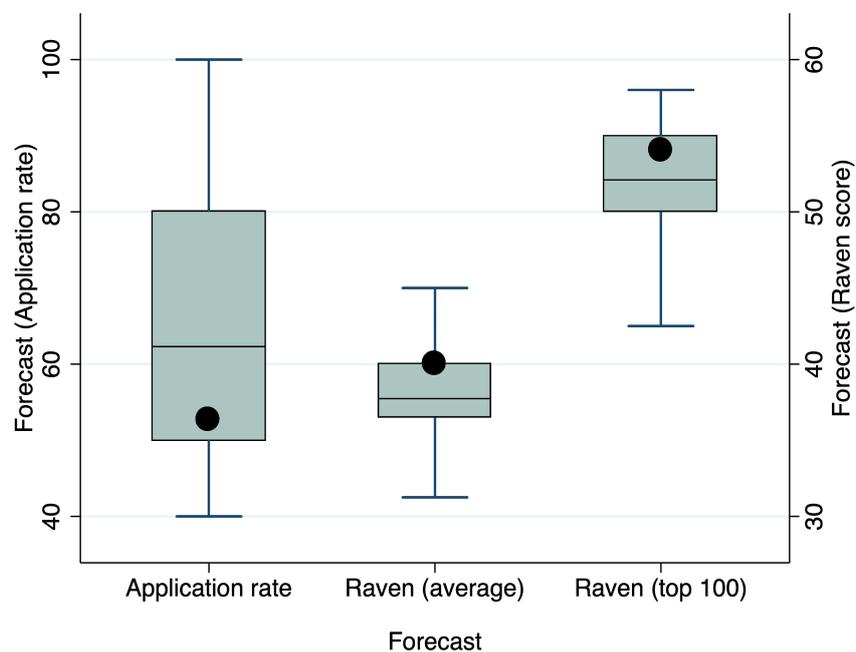
(b) Noisy ability, additional moments

(c) Noisy selection, core moments

(d) Noisy selection, additional moments

*Notes:* This figure plots the simulated moments against the empirical moments for all models presented in Table 4. All moments are normalised using the standard deviation of the respective empirical moment (which we obtain through a bootstrap). The underlying moments are reported in Table A48 -Table A51 of the online Appendix. For ease of representation, we plot the application rate in each treatment group (as opposed to the treatment effect on application rates) and average ability in each treatment group (as opposed to the treatment effect on average ability).

Figure 8: Forecast accuracy of firm managers



*Notes:* This figure reports managers' forecasts about the outcomes of the experiment. The circle shows the true value of the variable forecasted. The box-plot shows the distribution of the forecasts: (i) the horizontal line inside the box shows the mean forecast, (ii) the box shows the interquartile range of the forecasts, (iii) the upper and lower caps on the whiskers show the minimum and maximum forecasts. Sample used: all managers.

Table 1: Summary statistics and balance

	Mean		High wage (3)	StDev		N (5)	Balance tests ( $p$ )		
	Control (1)	Incentive (2)		Control (4)	Control (4)		Incentive = Wage = Control (6)	Incentive = Control (7)	Wage = Control (8)
Female	0.21	0.21	0.21	0.41	4689	0.98	0.88	0.88	
Age	26.08	25.95	26.24	4.28	4686	0.21	0.41	0.32	
Born in Addis Ababa	0.24	0.24	0.23	0.43	4689	0.53	0.72	0.45	
First language is Amharic	0.68	0.70	0.67	0.47	4689	0.21	0.15	0.89	
Heard about job in newspaper	0.55	0.58	0.56	0.50	4689	0.33	0.18	0.94	
Engineering or hard science	0.50	0.49	0.48	0.50	4689	0.46	0.35	0.24	
Economics	0.15	0.16	0.17	0.36	4689	0.53	0.38	0.29	
Other social science	0.15	0.16	0.14	0.36	4689	0.15	0.50	0.22	
Wage work experience (dummy)	0.53	0.53	0.53	0.50	4689	0.97	0.89	0.82	
Wage work experience (months)	28.12	28.45	29.06	43.53	4689	0.84	0.84	0.56	
Self-employed experience	0.33	0.35	0.35	0.47	4689	0.59	0.38	0.37	
Currently unemployed	0.67	0.65	0.64	0.47	4689	0.18	0.22	0.07	
Currently wage employed	0.24	0.26	0.27	0.43	4689	0.18	0.28	0.07	
GPA	2.97	2.99	2.96	0.46	4451	0.22	0.18	0.79	
Overall balance						0.33	0.61	0.47	

*Notes:* Balance tests. In columns 1-5, we report descriptive statistics and sample sizes. In columns 6-8, we report the  $p$ -values of a battery of balance tests. In column 6, we test the null hypothesis that the three experimental groups are balanced. In columns 7 and 8, we present pair-wise tests of a single treatment group against the control. In the last row, we report a joint test of orthogonality (following the recent literature, e.g. [McKenzie \(2017b\)](#)). To perform the joint test of orthogonality we regress the treatment variable on all covariates and we then test the joint null hypothesis that all covariates have a zero coefficient. We do this by estimating a categorical logit model in column 6 (the treatment variable can take three values) and OLS models in columns 7 and 8 (here we drop one experimental group, so the treatment variable only takes two values). Sample used: baseline sample.

Table 2: Application rates

	(1)
Incentive	0.115 (0.018)
High Wage	0.186 (0.017)
Control mean	0.412
Incentive = Wage ( $p$ )	0.000
Obs.	4689

*Notes:* OLS regression. The dependent variable is a dummy capturing whether the respondent has applied to the experiment's job. The second to last row reports the  $p$ -value of a test of the null hypothesis that the two treatments have the same effect. Robust standard errors reported in parenthesis. Sample used: baseline sample.

Table 3: Cognitive ability

	Mean	Percentile				
	(1)	90th (2)	75th (3)	50th (4)	25th (5)	10th (6)
Incentive	0.248 (0.112) [0.081]	0.229 (0.110) [0.115]	0.229 (0.117) [0.148]	0.170 (0.133) [0.607]	0.412 (0.173) [0.053]	0.079 (0.250) [1.000]
High Wage	0.194 (0.110) [0.225]	0.202 (0.108) [0.182]	0.227 (0.112) [0.130]	0.075 (0.131) [0.852]	0.280 (0.165) [0.271]	0.155 (0.227) [0.743]
Control value	-0.0000	2.312	1.477	0.356	-1.238	-2.697
Incentive = Wage ( $p$ )	0.574	0.795	0.983	0.448	0.371	0.741
Obs.	2386	2386	2386	2386	2386	2386

*Notes:* Estimates from OLS (Column 1) and quantile (Columns 2-6) regressions. The dependent variable is the index of cognitive ability. The second-to-last row reports the  $p$ -value of a test of the null hypothesis that the treatments have the same effect. Robust standard errors are reported in parenthesis. Sharp-ened  $q$ -values (Benjamini et al., 2006) are reported in brackets.  $q$ -values control the false discovery rate for the multiple tests of the same hypothesis for different indices of ability. A Wilcoxon rank-sum test rejects the equality of the distribution of cognitive ability in the control and incentive groups ( $p=.038$ ) and marginally fails to reject the equality of the distribution of cognitive ability in the control and wage groups ( $p=.107$ ). Sample used: all applicants.

Table 4: Parameter estimates

	(1)	(2)	(3)	(4)
Low $B$				
$\mu_C$	136.320 (20.69)	187.950 (15.99)	273.300 (32.21)	206.800 (234.69)
$\sigma_C$	192.840 (35.90)	192.010 (42.51)	206.450 (317.59)	208.570 (1839.90)
$\rho$	0.640 (0.06)	0.643 (0.05)	0.626 (0.21)	0.625 (0.22)
High $B$				
$\mu_C$	217.310 (15.66)	287.910 (14.53)	383.410 (52.01)	286.170 (60.86)
$\sigma_C$	245.100 (38.40)	239.910 (49.90)	260.400 (287.95)	263.320 (1182.06)
$\rho$	0.572 (0.05)	0.581 (0.05)	0.607 (0.23)	0.608 (0.25)
$a$	50.428 (1.19)	46.934 (0.82)		
$\mu_a$			-4.2e+03 (2.19)	792.460 (9.07)
$\sigma_a$			1.2e+04 (0.86)	1.4e+04 (0.69)
$\tau$	19.063 (15.93)	15.781 (13.80)	57.938 (91.39)	58.556 (358.42)
$\tau^w$	55.732 (53.50)	35.763 (87.08)	113.200 (306.60)	153.570 (1170.18)
Information Moments	Noisy ability 14	Noisy ability 16	Noisy selectivity 14	Noisy selectivity 16
Goodness of fit	2.2637	2.3389	2.2516	2.2654

Notes: The table shows parameter estimates obtained using minimum distance estimation. Columns (1) and (2) report estimates for the noisy-ability case. Columns (3) and (4) report estimates for the noisy-selectivity case. The estimates reported in columns (1) and (3) use fourteen empirical moments. The estimates reported in columns (2) and (4) use sixteen empirical moments. All empirical and simulated moments are reported in Tables A48 - A51. The parameters that describe the marginal distribution of ability are omitted for brevity and are reported in Table A47 in the online Appendix. Standard errors obtained through a bootstrap of the structural estimation reported in parenthesis. The bootstrap includes the estimation of  $B$  and the demediation procedure. Costs are expressed in Ethiopian Birr.

Table 5: The returns of the interventions

	Incentive given to...				High wage
	All applicants (1)	All hires (2)	Female (3)	Young (4)	(5)
Internal Rate of Return	9.8 (-18.7, 38.3)	382.5 (44.6, 720.4)	86.5 ( 4.3, 168.7)	47.1 (24.6, 69.6)	<0
	<i>Costs</i>				
Time costs (month 0)	386	133	77	207.9	485
Incentive (month 0)	5862	300	1017	2439.7	0
Wage costs (months 1-3)	0	0	0	0	4800
	<i>Benefits</i>				
Value of higher ability (months 1-45)	168	61	65	114.2	211

*Notes:* Cost-benefit analysis. We consider five interventions: (i) the application incentive as implemented in the experiment (column 1), (ii) an incentive offered conditional on meeting the selectivity threshold set by the employer in the application test (column 2), (iii) an incentive offered to all women (column 3), (iv) an incentive offered to all individuals below the median age (column 4), (v) the high wage intervention as implemented in the experiment (column 5). The figures reported in columns 1, 2 and 5 are based on the parameters estimated for the noisy-ability case, using the core set of moments. The first row reports the Internal Rate of Return (IRR) of each intervention. Below the IRR we report a 90 percent confidence interval obtained with a bootstrap. In the remaining rows we report costs and benefits expressed in Ethiopian Birr. The benefit of the intervention is given by the increase in the expected cognitive ability of hires, multiplied by the return to cognitive ability estimated using local labour market data. We consider three types of potential costs: (i) the time required to assess the additional applications, (ii) the cost of the incentive, (iii) the wage increase. All assumptions used in the computation of costs and benefits are discussed in detail in Appendix D.

Table 6: Firm managers' ranking of workers

	Ranked first		Ranked first or second
	(1)	(2)	(3)
Incentive	0.369 (0.053)	0.456 (0.064)	0.374 (0.055)
High wage	0.154 (0.048)		0.287 (0.062)
Control mean	0.159	0.272	0.446
Incentive = Wage ( $p$ )	0.001		0.069
No. managers	195	195	195
Obs.	585	390	585

*Notes:* OLS regression. The dependent variable is indicate in the column's heading. The unit of observation is an applicant-manager pair. We thus have three observations per manager. In column (2) we drop applicants from the high wage group. Standard errors clustered at the manager level reported in parenthesis. Sample used: all managers.