

DISCUSSION PAPER SERIES

IZA DP No. 16522

**Jobseekers' Beliefs about Comparative
Advantage and (Mis)Directed Search**

Andrea Kiss [®]
Robert Garlick [®]
Kate Orkin [®]
Lukas Hensel

OCTOBER 2023

DISCUSSION PAPER SERIES

IZA DP No. 16522

Jobseekers' Beliefs about Comparative Advantage and (Mis)Directed Search

Andrea Kiss ®

Carnegie Mellon University

Robert Garlick ®

Duke University

Kate Orkin ®

University of Oxford and IZA

Lukas Hensel

Peking University

OCTOBER 2023

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

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

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

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Jobseekers' Beliefs about Comparative Advantage and (Mis)Directed Search*

Worker sorting into tasks and occupations has long been recognized as an important feature of labor markets. But this sorting may be inefficient if jobseekers have inaccurate beliefs about their skills and therefore apply to jobs that do not match their skills. To test this idea, we measure young South African jobseekers' communication and numeracy skills and their beliefs about their skill levels. Many jobseekers believe they are better at the skill in which they score lower, relative to other jobseekers. These beliefs predict the skill requirements of jobs where they apply. In two field experiments, giving jobseekers their skill assessment results shifts their beliefs toward their assessment results. It also redirects their search toward jobs that value the skill in which they score relatively higher – using measures from administrative, incentivized task, and survey data – but does not increase total search effort. It also raises earnings and job quality, consistent with inefficient sorting due to limited information.

JEL Classification: J22, J24, J31, J64, O15

Keywords: search frictions, information frictions, occupational choice

Corresponding author:

Kate Orkin
University of Oxford
Wellington Square
Oxford OX1 2JD
United Kingdom
E-mail: kate.orkin@merton.ox.ac.uk

* Order of authors is randomized and all authors contributed equally. For helpful comments, we thank Johannes Abeler, Peter Arcidiacono, Linda Babcock, Vittorio Bassi, Michele Belot, Stefano Caria, Clement Imbert, David Huffman, Jeremy Magruder, Barbara Petrongolo, Chris Roth, Duncan Thomas, and Basit Zafar, as well as seminar and conference participants at AFE, AFES, Bocconi, CEPR, Carnegie Mellon, Columbia, Duke, Essex, FIU, IFLAME, NBER, Oxford, Peking, PHBS, Renmin, Pittsburgh, Royal Holloway, and SITE. We thank Alice Cahill, Raquel Caldeira, Aliya Chikte, Sabhya Gupta, Jenn Kades, Brynde Kreft, and Wim Louw for excellent research assistance. For collaboration on the broader research program, we thank Eliana Carranza, Maryana Iskander, Khodani Nematei, Laura Poswell, Neil Rankin, Mosoue Sekonyela, and Rob Urquhart. This study has been approved by the ethics review boards at the University of Cape Town (REC 2020/02/001), Duke University (# D0368), University of Oxford (# ECONCIA15-055), and Carnegie Mellon University (# STUDY2022_00000166). The experiments used in this study are preregistered on the AEA's trial registry at <https://www.socialscienceregistry.org/trials/1631/> and <https://www.socialscienceregistry.org/trials/10000/>. This paper has been produced through a research program financially supported by the World Bank Jobs Group, World Bank Africa Gender Innovation Lab, National Science Foundation (# 1824413), Private Enterprise Development in Low-Income Countries program (# 3024 and # 4728), UKRI GCRF Accelerating Achievement for Africa's Adolescents Hub, IPA Research Methods Initiative, and Upjohn Institute for Employment Research.

1 Introduction

Worker sorting into tasks and occupations has long been recognized as an important feature of labor markets (Roy, 1951). Efficiently matching workers with the tasks where their skills are most productive offers the prospect of large output gains (Lise & Postel-Vinay, 2020). The efficiency of matching depends crucially on job search, particularly how jobseekers direct their search effort across different job types. However, jobseekers may have imperfect information about how well their skills match what different job types require. Research in both developed and developing economies has documented imperfect information about aspects of job search but has not studied jobseekers’ beliefs about their skill match with different job types (Abebe et al., 2022; Abel et al., 2020; Bandiera et al., 2023; Beam, 2016; Belot et al., 2019; Boudreau et al., 2023; Conlon et al., 2018; Cortés et al., 2023; Jones & Santos, 2022; Kelley et al., 2023; Wiswall & Zafar, 2015).

This paper shows that jobseekers can have imperfect information about their skills, distorting how they direct job search and leading to worse labor market outcomes. We begin by presenting this idea in a simple model of directed job search. Jobseekers have multidimensional skills and beliefs about their **skill comparative advantage**: their ranks in different skill dimensions, relative to other jobseekers from similar backgrounds. These beliefs influence their **skill-directed job search**: how they allocate effort across searching for jobs with different skill demands. In this model, imperfect information about skill comparative advantage can distort skill-directed job search.

We provide causal evidence from two field experiments to support this idea. We work with a job search assistance agency in Johannesburg, South Africa. The agency helps young jobseekers from disadvantaged backgrounds, mostly with only secondary education, who are searching mainly for non-specialist, entry-level, service-sector jobs. We invite jobseekers to day-long job search assistance workshops run with the agency. In the workshops, we measure jobseekers’ skills using established psychometric assessments of communication and numeracy – two general-purpose, job-relevant skills that we verify firms in this context value in prospective workers. We also measure jobseekers’ beliefs about their levels of these skills relative to other jobseekers. Descriptive patterns in these data are consistent with our model and help to motivate our experimental analysis. We find that the difference between jobseekers’ beliefs about their skill comparative advantage and their assessment results is almost as large as if the beliefs were random guesses.¹

¹We define each jobseeker’s comparative advantage as the skill dimension in which they rank highest relative to other jobseekers from similar backgrounds who take the assessments. This follows the few job search studies using multidimensional skill measures (e.g. Guvenen et al. 2020) and common practice in education research (Altonji et al., 2016). We discuss trade-offs relative to other definitions later in the paper. Our theoretical and empirical arguments use this measure but the arguments do not rely on naming this measure “comparative advantage.” It could instead be named “higher-ranked skill” or some other term.

These beliefs persist over time and predict the types of jobs to which jobseekers apply.

In both experiments, we randomly assign half of the workshops to give participants information about their communication and numeracy assessment results, both expressed as rankings relative to a large group of jobseekers from similar backgrounds. We interpret the treatments as providing only jobseeker-facing information because it is difficult for participants to credibly share information with firms and we show that sharing is rare.

Our first experiment ($N = 278$) studies skill-directed job search using unusually rich survey, task, and administrative data on beliefs and search. Relative to the control group, treated participants' skill beliefs move substantially closer to their measured skills, including their beliefs about their comparative advantage. Treated participants are also more likely to apply to jobs whose skill demands match their comparative advantage, using multiple prespecified measures of skill-directed job search: applications on a job search platform, survey data, and a novel incentivized task in which participants choose between applying to jobs with different skill requirements.² These results are driven by jobseekers who get more information from treatment: those whose baseline comparative advantage belief *did not* match their assessment results. Treatment closes roughly 80% of the gap in skill-directed job search between these jobseekers and those jobseekers whose baseline comparative advantage belief *did* match their assessment results.

Our second experiment ($N = 4,389$) studies labor market outcomes using survey data for a larger sample collected on average 3.5 months after the workshops. Treated jobseekers have weekly earnings and hourly wages roughly 25% higher than control group jobseekers and are more likely to have formal jobs. But we find at most weak evidence for a higher employment rate, which may be a more difficult margin to move. Treatment effects on jobseekers' comparative advantage beliefs and skill-directed job search are qualitatively similar to the first experiment, although using different measures.

The two experiments show a consistent picture of this labor market: getting more information about their skills shifts jobseekers' perceived comparative advantage toward their measured comparative advantage, redirects their search toward jobs aligned with their measured comparative advantage, and leads to higher-quality employment. In contrast, we find limited evidence for a plausible alternative mechanism: that treatment shifts beliefs about absolute skill levels and therefore shifts search effort and labor market outcomes. We do see the first part of this mechanism in our data: treatment lowers the average jobseeker's belief about her skill level, because untreated jobseekers are on average overconfident about their skill level relative to other jobseekers. But we do not see the

²This task provides a novel way to use job application choices to measure jobseekers' valuation of skill match relative to other vacancy characteristics. This builds on work using choices in controlled environments to study preferences over education, jobs, and occupations ([Adams-Prassl & Andrew, 2023](#); [Mas & Pallais, 2017](#); [Wiswall & Zafar, 2015](#)).

second part of this mechanism: treatment has negligible effects on job search effort in both experiments, using multiple measures. This is consistent with our model, which shows that the sign of the effect of skill level beliefs on search effort is theoretically ambiguous.

The two experiments provide complementary types of evidence, from the same location, with participants recruited in the same way, several years apart. Given the two experiments' different strengths, we call them the "tight" and "big" experiments respectively. We use this light-hearted, nonstandard terminology because standard taxonomies do not capture the differences between the two experiments.³ Together, the two experiments give us exogenous variation in jobseekers' beliefs about their skills and the unique combination of data needed for this research: data on jobseekers' multidimensional skills, beliefs about these skills, and job search activities and outcomes, including data on their application decisions at specific vacancies and the skill demands of these vacancies.⁴ This adds to a small, recent literature showing the value of combining multiple research designs and data sources to understand both a job search process and the outcomes of that process (Carranza et al., 2022; Cortés et al., 2023; Field et al., 2023).

More generally, what labor market conditions might lead to inaccurate comparative advantage beliefs and hence to poorly skill-directed job search? Our sample enters the labor market with limited information about their comparative advantage because schools give noisy feedback on their skills (Lam et al., 2011). The same is true in many education systems (Pritchett, 2013). Our sample learns slowly about their comparative advantage because high unemployment limits learning about skills through work experience. Similarly, many countries face high and persistent youth unemployment (ILO, 2022). Misdirected search may be particularly important when search is costly for jobseekers or screening mismatched applicants is costly to firms (Abebe et al., 2021a; Algan et al., 2022; Fernando et al., 2023; Hensel et al., 2022).

This paper provides the first direct evidence that jobseekers' beliefs about their skill comparative advantage influence how they direct search over different job types and therefore influence their labor market outcomes. This extends and bridges two literatures: on belief-dependent job search and on skill (mis)match between jobs and workers.

First, we extend recent research on the relationships between jobseekers' beliefs, search activities, and search outcomes, reviewed by Mueller & Spinnewijn (2023).⁵ This litera-

³Both are field not lab or lab-in-the-field experiments because they feature real jobseekers engaging in real job search. Both have elements of framed and natural field experiments (Harrison & List, 2004).

⁴No existing data sources provide all this information. Survey datasets typically record aggregated search measures such as total applications submitted, not jobseeker $\times$ vacancy-level measures. Job search platforms seldom measure skills, beliefs, or labor market outcomes. Government administrative data do not measure beliefs and seldom measure skills or search.

⁵This includes research into the co-evolution of search and search-related beliefs in panel data (Adams-Prassl et al., 2023; Conlon et al., 2018; Mueller et al., 2021; Spinnewijn, 2015); experiments providing in-

ture focuses on jobseekers’ beliefs about the level of their labor market prospects, captured by job offer arrival rates or wage offer distributions. Despite the centrality of sorting on skill comparative advantage in labor economics (Roy, 1951), the literature on belief-based job search has not studied beliefs about skills or skill comparative advantage. And it rarely studies how search is directed across different job types and what this implies for job-worker matching. The most similar work shows that encouraging jobseekers to apply to different occupations can redirect job search and improve some labor market outcomes (Altmann et al., 2022; Belot et al., 2019, 2022). We complement this work by focusing on skills rather than occupations and directly measuring the beliefs that link these treatments to search decisions. We also show how to directly measure the entire causal chain – skills, beliefs about skills, skill-directed job search, and search outcomes – when working in a setting without government administrative data.

Second, our work contributes to a broad and long-standing literature studying barriers to efficient matching between workers and jobs including capital market frictions (Banerjee & Newman, 1993), discrimination (Goldin, 1990), migration costs (Bardhan & Udry, 1999), and limited information (Jovanovic, 1979). One part of this literature emphasizes the importance of the (mis)match between workers’ multidimensional skills and jobs’ skill demand (Sanders & Taber, 2012) but typically relies on indirect evidence for the search process generating mismatch. In particular, Baley et al. (2022), Fredriksson et al. (2018), Guvenen et al. (2020) and Lise & Postel-Vinay (2020) show that long-term earnings losses may arise when jobseekers have imperfect information about how their multidimensional skills match with different jobs. These papers do not observe beliefs or search, instead inferring these processes from dynamic search and matching models fitted to rich longitudinal employment data. Our measures of skills, beliefs about skills, search, and short-term labor market outcomes, combined with experimental variation in beliefs, allow us to directly observe behavior assumed by these models. Within this literature, our findings are most consistent with the modeling framework of Baley et al. (2022), which features jobseekers directing search based on inaccurate beliefs about their multidimensional skills. The limited information and misdirected search we document might also help to explain aggregate variation in job turnover rates (Donovan et al., 2023).

formation about labor market conditions (Altmann et al., 2018; Jäger et al., 2023; Jones & Santos, 2022); experiments on search subsidies, matching services, or mentoring programs that influence multiple outcomes including jobseekers’ beliefs (Abebe et al., 2022; Alfonsi et al., 2022; Bandiera et al., 2023; Banerjee & Sequeira, 2023; Beam, 2016; Kelley et al., 2023; Wheeler et al., 2022); and research on jobseekers’ beliefs about attributes of specific jobs (e.g. Bazzi et al. 2021; Boudreau et al. 2023; Chakravorty et al. 2023; Sockin & Sojourner 2023; Subramanian 2022). The formation and consequences of beliefs about skills have also been studied outside of job search: decisions in the workplace (e.g. Hoffman & Burks 2020; Huffman et al. 2022; Malmendier & Tate 2015), education investment (e.g. Berry et al. 2022; Bobba & Frisncho 2022; Dizon-Ross 2019), and in lab settings, including some studies of tasks that mimic job search (e.g. Falk et al. 2006). Skill comparative advantage beliefs have been studied in education but not in other settings (Altonji et al., 2016).

A related literature focuses on firm-side limited information about jobseekers’ skills.⁶ Within this literature, our work is most closely related to our companion paper, [Carranza et al. \(2022\)](#). That paper shows that new information about jobseekers’ skills has larger labor market effects when revealed to firms and jobseekers than when revealed only to jobseekers. It uses multiple experiments to make that argument, including reporting the labor market outcomes from our big experiment. But it only briefly explores jobseeker-side mechanisms for these labor market outcomes. In this new paper, we focus deeply on jobseeker beliefs and behavior by adding the tight experiment, with its detailed measures of skill-directed job search and comparative advantage beliefs. We also add new mechanism measures from the big experiment and a model of belief-based skill-directed job search. Combining the tight and big experiments allows us to study the full model-guided causal chain from comparative advantage beliefs, through skill-directed job search, to labor market outcomes, which the companion paper does not do.

Our findings also relate to extensive work on active labor market programs in both developed and developing economies, reviewed by [Card et al. \(2018\)](#) and [McKenzie \(2017\)](#). Our skill information intervention raises the average treated participant’s earnings by more than enough to cover the average variable cost of the intervention. This suggests similar interventions might be cost-effective additions to active labor market programs.

Section 2 of the paper describes our model, context, and patterns of skill, skill beliefs, and job search in our sample. In Section 3 we show the relationship between comparative advantage beliefs and skill-directed job search in the tight experiment. In Section 4 we show the relationship between comparative advantage beliefs, skill-directed job search, and labor market outcomes in the big experiment. We show that treatment has little effect on search effort in Section 5 and on other possible mechanisms in Section 6. We informally discuss possible general equilibrium considerations in Section 7.

2 Economic Environment

We begin with a conceptual framework and then describe our context, sample, and skill assessments. We then report four key descriptive patterns that inform our conceptual framework: jobseekers’ skills vary across multiple dimensions, different firms value different skill dimensions, jobseekers’ beliefs about their skills persistently differ from results on assessments, and jobseekers’ beliefs about their skills predict their search decisions. These four patterns motivate the structure of our conceptual framework, which shows

⁶Hiring and wage-setting can change when firms observe new information about workers’ skills from job performance ([Altonji & Pierret, 2001](#); [Arcidiacono et al., 2010](#); [Hardy & McCasland, 2023](#); [Kahn & Lange, 2014](#)), references and referrals ([Abel et al., 2020](#); [Heath, 2018](#); [Ioannides & Loury, 2004](#); [Pallais, 2014](#)), education qualifications ([Alfonsi et al., 2020](#); [Clark & Martorell, 2014](#); [Jepsen et al., 2016](#); [MacLeod et al., 2017](#)), and skill certification ([Abebe et al., 2021b](#); [Bassi & Nansamba, 2022](#); [Groh et al., 2015](#)).

how jobseekers' limited information about their comparative advantage over multiple skills can distort skill-directed job search and worsen labor market outcomes.

In the conceptual framework and tight experiment, we study jobseekers' beliefs about comparative advantage over communication and numeracy and their search over jobs demanding these skills. Two skill dimensions allow simplicity and are the minimum needed to study comparative advantage and skill-directed search. These skills suit our research question: they are general-purpose skills used in many jobs in this economy, they are weakly correlated with each other, and different firms value them differently. The big experiment shows that results generalize to a less stylized setting where jobseekers receive information about six skills and search over jobs requiring many skills.

2.1 Conceptual Framework

Our conceptual framework has a similar spirit to recent models of “partially directed job search,” where jobseekers try to direct search to higher-wage vacancies but face uncertainty about wages (Lentz et al., 2022; Wu, 2021). We use a static partial equilibrium framework focusing on jobseekers' search and beliefs, treating labor demand and wage posting as fixed, because we observe limited firm-side data.

We assume that each jobseeker has communication and numeracy skill levels S_C and S_N . Each job demands primarily communication or primarily numeracy skills.

Search over jobs demanding different skills: Jobseekers split fixed total search effort $\bar{E}$ between search for communication jobs E_C and numeracy jobs E_N . We generalize the framework later to allow endogenous choice of total effort. Searching for a type j job yields outcome $V_j(S_C, S_N, E_j)$, which is a reduced-form expression for the expected present value of a job offer scaled by the probability of an offer. We make three assumptions. First, V_j is increasing and concave in all three arguments and $\partial V_j / \partial S_j > \partial V_j / \partial S_i > 0$ for $j \neq i$. This assumption allows both types of jobs to value both skills, but each job type to value one skill more. Second, we assume that skill and search effort are technical complements and are ‘more complementary’ within than across dimensions. Intuitively, a jobseeker with high communication skills will get a higher return to directing marginal search effort to communication than numeracy jobs and vice versa. Formally,

$$\frac{\partial^2 V_j}{\partial S_j \partial E_j} > \frac{\partial^2 V_i}{\partial S_j \partial E_i} > 0 \quad (1)$$

for $j \neq i$. Third, gross utility from job search $U(V_C, V_N)$ is increasing and concave in both arguments. This allows jobseekers to value the outcomes of searching for both job types without fully specifying the offer acceptance decision or the reservation wage.

Under the first and third assumptions, jobseekers direct search effort to equalize the

marginal utility of searching for each job type:

$$\frac{\partial U}{\partial V_C} \times \frac{\partial V_C}{\partial E_C} = \frac{\partial U}{\partial V_N} \times \frac{\partial V_N}{\partial E_N}, \quad (2)$$

where $\frac{\partial U}{\partial V_j}$ captures the jobseeker's preferences over nonpecuniary aspects of job type j . Conditional on these preferences, marginal search effort will be directed based on the relative magnitudes of $\frac{\partial V_C}{\partial E_C}$ and $\frac{\partial V_N}{\partial E_N}$. Under the second assumption, $\frac{\partial V_C}{\partial E_C}$ is more steeply increasing in communication skill than $\frac{\partial V_N}{\partial E_N}$. This means that if a jobseeker's communication skill rises, the left-hand side of the optimality condition in (2) will rise more than the right-hand side. The jobseeker will then increase E_C and decrease E_N to restore equality in condition (2), under the assumption that V_j is a concave function of E_j .

Jobseekers' beliefs about their skills: We assume each jobseeker has beliefs about their skill levels $\tilde{S}_C$ and $\tilde{S}_N$. They allocate search effort based on these beliefs, not their actual skill levels. We do not model belief formation, including belief updating in response to search outcomes, because we show in Section 2.6 the jobseekers we study learn little about their skills during search over the timeframe of our study.

Testable predictions: Our framework predicts two effects of a jobseeker receiving accurate new information about her skill levels, for example, learning that her communication skill is higher than she previously thought. First, she will redirect search effort away from numeracy-heavy jobs and toward communication-heavy jobs, because condition (1) means that her expected relative return to search is now higher for communication- than numeracy-heavy jobs. Second, her labor market outcomes will improve, because more of her search effort is now directed to job types that will reward her skills more. Note that these outcomes are driven by learning about her communication skill relative to numeracy skill. Learning about her average skill level across the two dimensions is irrelevant in this framework because total search effort is fixed. We later generalize the framework to allow her belief about skill level to endogenously influence total search effort but show empirically that this does not occur in our data.

Skill comparative advantage: We define a jobseeker as having a comparative advantage in communication if she ranks higher in the distribution of communication than numeracy skills and vice versa. This definition aligns with the spirit of standard definitions of comparative advantage in trade. There, a country has a comparative advantage in the product it can produce at lowest opportunity cost. Here, a jobseeker has a comparative advantage in the skill where she ranks higher, because she can supply it to the market at a lower opportunity cost in terms of time spent supplying the other skill.

This definition aligns with our empirical work, where we rank each jobseeker's communication and numeracy skills using her scores on standardized assessments relative to

a large group of jobseekers from similar backgrounds. We focus on relative ranks in each skill distribution rather than absolute scores on assessments because ranks are familiar outcomes from many education systems and because absolute scores are sensitive to the way assessments are scaled (Nielsen, 2023). While other definitions are possible, this definition takes advantage of our multidimensional skill measures and follows applications when researchers have access to similar data in education (Altonji et al. 2016) and in labor research (e.g. Guvenen et al. 2020).⁷ Our empirical analysis uses jobseekers’ relative ranks in the communication versus numeracy skill distributions. But neither our theoretical predictions nor our empirical analysis relies on naming this measure “skill comparative advantage” rather than “relative skill rank” or some other term.

Using this definition of comparative advantage, our framework’s testable predictions become that a jobseeker who gets accurate new information about her skill comparative advantage will (1) redirect search effort toward jobs whose skill demand aligns with her comparative advantage and (2) have better labor market outcomes.

2.2 Context and Target Population

We work in Johannesburg, part of South Africa’s commercial and industrial hub of Gauteng, a metropolitan area of 14 million people. Wage labor is the primary income source, self-employment is low, entry-level employment is mainly in services and manufacturing, and most employment is in formal firms (SSA, 2022). This matches patterns in large cities in other middle-income African countries (Bandiera et al., 2022a).

Our target population is young, active jobseekers with at least high school education who attended school in low-income areas. We recruit from participant registries at the Harambee Youth Employment Accelerator. Harambee is a social enterprise funded jointly by the South African government and private firms that provides job search assistance to young jobseekers from low-income backgrounds. Since 2013, Harambee has maintained a database of active jobseekers recruited through traditional and social media. Jobseekers sign up online with their national identity number, which Harambee uses

⁷Because we rank each jobseeker’s skills relative to those of other jobseekers from similar backgrounds, our definition is designed to approximate each jobseeker’s comparative advantage relative to likely competitors for similar jobs. However, we do not observe the jobseeker’s rank in the skill distribution of applicants for any specific job. Alternative approaches estimate occupation-specific wages by education and use this to define comparative advantage in occupations based on education level (e.g. Acemoglu & Autor 2011; Gibbons et al. 2005). These approaches can price the value of different types of workers in different types of jobs, which our approach does not. But their wage estimates are conditional on the way workers currently sort into occupations. These may be sensitive to a Lucas-style critique that wages might be different under a different type of sorting, which is precisely the mechanism that we study. We show in Section 3.4 that demands for communication and numeracy skills are roughly equal in this setting, using a very rough proxy of skill demand. This suggests that defining comparative advantage using relative skill ranks and using skill prices might not produce completely different classifications in this setting.

to determine that they are aged 18–34 with legal permission to work. They self-report if they are actively searching for work and attended school in a low-income area. Firms receive free access to the database for recruiting. The database captures a sizeable proportion of the population of interest: in Gauteng in 2022, restricting to ages 18–34, there were 1,078,745 jobseekers in the database and 1,403,064 unemployed people (SSA, 2022). These groups do not perfectly overlap because some jobseekers in the database are employed and some unemployed people do not sign up on the database.

This context and population are relevant for studying limited information about skills, although they are not nationally or globally representative. These jobseekers are actively engaged in the labor market but face an economic environment that provides limited information about their skill levels relative to other jobseekers, for two reasons.

First, jobseekers from low-income schools receive limited information about their skills from assessments during school. High school grades and grade progression are only weakly correlated with results on independent skill assessments (Lam et al., 2011). Such schools rarely have career counselors and few teachers are trained to offer career guidance (Pillay, 2020). There is a national high school graduation exam but exam grades are not strongly informative: they weakly correlate with performance in post-secondary education and firms report in interviews that the grades convey limited information about skills (Schoer et al., 2010). Second, unemployment is high: the unemployment rate for ages 15–34 in Johannesburg at the time of the tight experiment was 40.5% (SSA, 2022).⁸ Many jobseekers have no work experience several years after completing education. This limits their scope to learn about their skills through work experience, the main mechanism in models of worker-side learning in other contexts (Baley et al., 2022). These patterns are not unique to South Africa, as we explain on page 4.

2.3 Sample Description & Skill Assessments

Recruitment: We recruited one sample of 4389 people for our big experiment between September 2016 and April 2017 and one sample of 278 people between July and October 2022 for our tight experiment, using the same sampling frame and method. We contacted people from Harambee’s database who lived within commuting distance of our field location in downtown Johannesburg. We screened out those not actively searching for work and invited the rest to a day-long job search assistance workshop. We stated that the workshop would include taking assessments that could be used to match them to suitable vacancies, answering questions about their job search, and receiving job search

⁸We use Statistics South Africa’s definition of an employed person: someone who engaged in any income-generating activity, for at least one hour, during the reference week. Unemployment rates exclude those in full-time education or not in the labor force.

advice. Harambee often runs similar workshops.

Data: In both experiments, participants complete self-administered baseline surveys about their demographics, beliefs about their skills and labor market prospects, recent job search activities, current employment, and employment history. They also take skill assessments in person. We discuss post-treatment measurement in sections 3 and 4.

For the tight experiment, we also observe participants’ search behavior on the online job search and matching platform [SAYouth.mobi](#). Harambee used its database of participants to set up this platform in 2019. The platform aggregates job advertisements from all online job boards that jobseekers can access without incurring mobile phone data charges. Jobseekers in our sample described the platform as a key part of their job search strategy. It also allows firms to post advertisements and set up interviews.

Sample characteristics: We focus here on descriptive statistics for the tight experiment sample, shown in Table 1. We describe the big experiment sample in Section 4.1. Jobseekers are young, with interdecile age range 21–32. Most, 60%, have only secondary education, where there is little specialization by subject, limiting their scope to learn about their skills from specialized training. Another 37% have some post-secondary education. They have limited, mostly informal work experience so have had limited opportunity to learn about their skills from work: 33% were employed at baseline but only 13% had a formal written contract and only 25% had *ever* held a long-term wage job.

Their job search effort was high but met with limited success. 96% were actively searching. In the week before baseline, the average jobseeker submitted 10 job applications, spent 14 hours searching, and spent 22.72 USD on search, including transport to drop off CVs and attend interviews, mobile phone airtime and data, and printing.⁹ This search cost is high: 50% of the average weekly earnings of employed people in our sample. This matches other work on job search costs in South Africa ([Banerjee & Sequeira, 2023](#); [Kerr, 2017](#)). High search costs limit scope for jobseekers to learn about their skills through search and raise costs of misdirected job search. The average jobseeker received only 0.17 job offers in the preceding 30 days, implying that < 1% of applications yield offers. (We used a longer recall period for offers than applications because offers are rare.)

Skill assessments: The numeracy assessment captures practical arithmetic and pattern recognition. It was developed by a large retail chain to assess potential cashiers. The communication assessment captures English-language listening, reading, and comprehension skills at a high school level. It was developed by an adult education provider ([www.mediaworks.co.za](#)). Candidates also complete a “concept formation” assessment that captures fluid intelligence or ability to see identify patterns across situations and to

⁹All monetary values throughout the paper are in 2021 USD in purchasing power parity terms.

Table 1: Summary Statistics – Tight Experiment

	Mean (1)	Median (2)	Min (3)	Max (4)	SD (5)	Obs. (6)
<u>Panel A: Demographics</u>						
Black African	1.00	1.00	1.00	1.00	0.00	278
Male	0.33	0.00	0.00	1.00	0.47	278
Age	26.41	26.00	18.00	36.00	4.04	278
Completed secondary education only	0.60	1.00	0.00	1.00	0.49	278
University degree / diploma	0.22	0.00	0.00	1.00	0.41	278
Any other post-secondary qualification	0.15	0.00	0.00	1.00	0.36	278
<u>Panel B: Labor market background</u>						
Any work in last 7 days	0.33	0.00	0.00	1.00	0.47	278
Has worked in permanent wage job before	0.25	0.00	0.00	1.00	0.43	278
Earnings in USD (last 7 days, winsorized)	44.90	0.00	0.00	697.57	101.71	277
Written contract	0.13	0.00	0.00	1.00	0.33	278
<u>Panel C: Search behavior</u>						
Any job search in last 30 days	0.96	1.00	0.00	1.00	0.20	278
# applications (last 7 days, winsorized)	10.00	5.00	0.00	100.00	14.93	278
Search expenditure in USD (last 7 days, winsorized)	22.72	14.00	0.00	126.00	23.72	278
Hours spent searching (last 7 days, winsorized)	13.82	9.00	0.00	72.00	15.00	278
# job offers (last 30 days, winsorized)	0.17	0.00	0.00	4.00	0.56	278
<u>Panel D: Skills beliefs</u>						
Aligned belief about CA	0.49	0.00	0.00	1.00	0.50	278
Fraction aligned belief domains	0.22	0.00	0.00	1.00	0.31	278

Notes: Table 1 shows baseline summary statistics for the tight experiment. ‘CA’ stands for comparative advantage. Winsorization is at the 99th percentile. All monetary values are in 2021 USD PPP.

use logic in new situations (Raven & Raven, 2003; Taylor, 1994). These assessments are also used in the big experiment, along with three more assessments described in Section 4.2. We describe the assessment process and psychometric properties of the assessments in Appendix C. Assessment results show no ceiling or floor effects (Figure B.1).

All our information treatments and measures of skill beliefs use assessment results relative to a benchmark population: roughly 12,000 jobseekers from Harambee’s database who have been assessed. We place jobseekers’ communication and numeracy skills in quintiles relative to the distribution of assessment scores in this population.

2.4 Skills Vary Substantially across Dimensions within Jobseeker

These two assessments differentiate jobseekers horizontally more than vertically: they show which jobseekers are better suited for jobs where communication versus numeracy skills are more important, rather than identifying jobseekers who are likely to be better at most jobs. In the full tight experiment sample, communication and numeracy scores have correlation 0.31 and only 17% of jobseekers are in the top two quintiles for both

skills. Communication and numeracy scores are both moderately correlated with concept formation ($\rho \approx 0.25$ for both skills), suggesting neither assessment captures ability better.

In the tight experiment, we restrict our sample to 278 jobseekers with a unique skill comparative advantage. This means that we exclude 95 of the 373 jobseekers who participate in the experiment because they score in the same quintile for communication as for numeracy. We impose this restriction because our experimental measures of skill-directed job search, described in Section 3.4, can be sensibly defined only for jobseekers with a unique skill comparative advantage. The shares of this sample with comparative advantage in communication and numeracy are respectively 62 and 38% (Table B.1).¹⁰

2.5 Firms Value Different Skills for Different Jobs

We have already shown that the assessments mainly identify jobseekers' relative strengths across different skills. Here we show that firms value the communication and numeracy skills that we measure and that there is variation in which firms value which skills. These patterns suggest scope for jobseekers to improve their labor market outcomes by searching for jobs that value their skills on which they score highly.

Firms value communication and numeracy skills relative to education qualifications: In an incentivized resume-ranking experiment, 91% of firms preferred job applicants with higher communication or numeracy assessment results to job applicants with lower assessment results and an additional one-year post-secondary training qualification. (See Appendix B.2 for details.) Over 500 client firms have paid Harambee to screen roughly 1 million jobseekers using these assessments, which we interpret as revealing a preference for using these skills in hiring. And Carranza et al. (2022) use multiple field experiments to show that jobseekers who can certify their assessment results have better labor market outcomes. We do not argue that communication and numeracy are the most important skills or only relevant skills in this labor market, just that they are important.

These results are consistent with the national high school graduation exam not giving firms sufficient information on applicants' skills. The communication and numeracy skill quintiles are positively but weakly associated with jobseekers' self-reported grades on their graduation exams in English and mathematics, respectively (Table C.3, columns 1–2). The positive association suggests the Harambee assessments capture meaningful variation in skills; the weak association highlights that jobseekers and firms both have scope to learn from information on jobseekers' ranks on Harambee's assessments.

Different firms value different skills: In the same incentivized ranking experiment, 58% of the firms in our sample ranked a resume with high numeracy skills ahead of a

¹⁰The shares of jobseekers in this sample with communication and numeracy comparative advantage do not need to be equal because we use a separate benchmark population.

resume with high communication skills, and 42% of firms had the opposite ranking. This cross-firm variation, combined with the fact that our assessments horizontally differentiate jobseekers on communication versus numeracy skills, creates scope for jobseekers to improve their labor market outcomes by searching for different types of jobs.

Firms can at least partly observe skills. We conducted a measurement exercise embedded in a firm’s hiring process, described in Appendix B.3. We find that the firm’s HR team can identify jobseekers’ skill levels based on typical application materials. The HR team is also significantly more likely to invite applicants to interviews for positions with skill requirements that match jobseekers’ assessed comparative advantage, than to interviews for positions that do not match their comparative advantage. Together, these patterns suggest that redirecting jobseekers’ search towards jobs that match their comparative advantage in skills has the potential to improve their labor market outcomes.

Our companion paper uses another experiment to show that firms do not fully observe jobseekers’ skills from application materials, and application outcomes change when firms also observe skill assessment results (Carranza et al., 2022). That finding is not inconsistent with the evidence we discuss in this subsection. Together, they show that firms can partly but not perfectly observe jobseekers’ skills.

2.6 Jobseekers’ Perceived and Measured Comparative Advantage in Skills Differ

We have already shown scope for skill-directed job search to improve jobseekers’ labor market outcomes. Here we show that jobseekers’ beliefs about their skills do not match their skill assessment results, which might distort search direction.

We measure jobseekers’ beliefs about their communication and numeracy skill quintiles before they take assessments. We first define the skills, then explain the concept of quintiles, define the reference group, and ask jobseekers which quintile they are in on each skill, relative to the reference group. See Appendix C.1 for measurement details and Figure D.1 for the sequencing of belief measurement, assessments, and treatment. We do not refer to jobseekers’ beliefs about skills as “accurate” or “inaccurate,” as no assessment perfectly measures skills, even the well-established assessments we use. Instead, we refer to beliefs about skills as “aligned” or “misaligned” with skill assessment results.

The joint distribution of jobseekers’ skills and beliefs is a multidimensional object that can be described in many different ways. We focus here on two simple, binary measures of belief (mis)alignment that are guided by our conceptual framework. We show robustness checks later using other, continuous measures of belief (mis)alignment.

Jobseekers have misaligned beliefs about their comparative advantage between skills: We define a jobseeker as having an **aligned comparative advantage belief** about

her comparative advantage over communication and numeracy if she believes she is a higher quintile for the same skill in which she actually scores in a higher quintile on our assessments. Using this definition, only 49% of jobseekers have aligned comparative advantage beliefs at baseline (Table 1, panel D).¹¹

The measure above captures jobseekers' beliefs about their communication and numeracy skills, not their assessment results. We also ask jobseekers' which quintile they fall in on our communication and numeracy assessments, after taking the assessments but before any jobseekers receive information about their skills. These two belief measures may differ if, for example, a jobseeker believes she has a comparative advantage in communication but that she happened to score badly on our specific communication assessment. But in practice the two beliefs measures are highly correlated and show similar levels of misalignment with assessment results.¹² Only 54% of jobseekers' comparative advantage beliefs about their assessment results match the actual results.

Jobseekers have misaligned beliefs about their average level over the two skills: We define the **fraction of aligned beliefs** as the average of two indicators, one for communication and one for numeracy, each equal to one if the perceived and assessed skill quintiles are equal. The measure can take values of 0, 0.5 and 1. Using this measure, 22% of beliefs about skills are aligned (Table 1, panel D).¹³ The misalignment is explained more by overconfidence than underconfidence: 63% of jobseekers' beliefs about their skills are above their assessed skills and only 19% are below. Similar patterns of overconfidence have been documented in other work (Santos-Pinto & de la Rosa, 2020).¹⁴

Jobseekers learn little about their skills while searching: To show this, we use the big experiment's baseline and follow-up surveys. The share of control group jobseekers whose believed and assessed comparative advantage align hardly changes over 3.5 months, even for the employed and those with above-median search effort (Table C.5,

¹¹If jobseekers were randomly guessing, 40% would have aligned comparative advantage beliefs. To see this, note that there are 5x5 possible beliefs over communication and numeracy skill quintiles. The 5 possible beliefs with tied skill quintiles are misaligned because our sample excludes anyone with tied assessment results. 10 of the other 20 possible skill beliefs lead to misaligned comparative advantage beliefs. Hence random guessing leads to aligned comparative advantage beliefs in $(5 \times 0 + 20 \times 0.5) / 25 = 40\%$ of cases.

¹²The high correlation suggests that jobseekers view the assessments as capturing relevant parts of their general skills. Regressing general beliefs about skills on beliefs about assessment results for the same skill produces coefficients of 0.39–0.52, controlling for assessment results, demographics, and education.

¹³Comparative advantage beliefs and the fraction of aligned beliefs are positively correlated by construction but have some separate variation ($\rho = 0.21$). To see the separate variation, consider a candidate who has measured skills in quintile 2 in communication and 4 in numeracy, and has aligned beliefs. She scores 1 on both belief measures. Raising her believed numeracy quintile without changing her communication belief will decrease the fraction of aligned beliefs to 0.5 without changing aligned comparative advantage.

¹⁴Some researchers call this pattern, in which people think they are performing better than others, “overplacement” (Moore & Healy, 2008). We call this pattern overconfidence, except where otherwise noted.

columns 1–3). The fraction of aligned beliefs is similarly persistent (columns 4–6). Slow learning might reflect limited scope for feedback during search: only 3% of jobseekers report ever receiving feedback about their skills during an unsuccessful job application. It may also reflect the well-documented difficulty of Bayesian learning about a multi-input function, such as the search outcome-skill relationship (Banerjee & Sequeira, 2023).

Jobseekers seem to draw on both school results and other information in forming skill beliefs: Even with good information from the schooling system, high school graduation exam scores should not perfectly predict young adults’ beliefs about their skills: many jobseekers took school exams multiple years ago, the exam and assessments do not test identical skills, and jobseekers may not perfectly recall their exam scores. Indeed, jobseekers’ beliefs about their communication and numeracy skills are positively but weakly correlated with their self-reported results on the graduation exams in English and mathematics (Table C.3, columns 4–5). Beliefs about comparative advantage are positively associated with the difference in scores between the two exam subjects (columns 6–7).

Skill beliefs do not substantially vary by gender: In both the tight and big experiments, we find limited gender differences in baseline skill beliefs, with or without controls for assessment results, demographics, and education (Appendix J, Tables J.1 and J.2). This matches a recent metastudy showing limited gender heterogeneity in confidence (Bandiera et al., 2022b). This motivates our gender-pooled analysis in this section.

2.7 (Misaligned) Beliefs Predict Skill-Directed Job Search

We have already shown that there is scope for skill-directed job search to improve jobseekers’ labor market outcomes but that jobseekers’ beliefs about their skills don’t match their assessment results. Here we show that the gap between jobseekers’ beliefs about their skills and assessment results might shift their skill-directed job search.

Comparative advantage beliefs predict job search decisions: In the big experiment, we ask candidates what skill is most valuable for the types of jobs for which they are applying. Candidates’ answers are strongly associated with their beliefs about their comparative advantage in skills in the control group. They are 9–10pp ($p < 0.01$) more likely to state that they are applying for jobs that value the skill in which they *believe* they have a comparative advantage, compared to jobs valuing other skills (Table C.6, rows 1–2). The correlations are robust across skill domains and to controlling for measured comparative advantage and demographics. These results suggest jobseekers try to search for jobs for which they have a skill comparative advantage.

Assessed comparative advantage weakly predicts job search decisions: Jobseekers are only 2–5 pp more likely to state that they are applying for jobs that value the skill in

which they *scored higher on our assessments*, compared to jobs valuing other skills (Table C.6, rows 3–4). The results in this and the preceding paragraph suggest that jobseekers’ skill beliefs might direct their search away from jobs that match their assessed skills.

Jobseekers also believe that having higher skills improves search outcomes. In the big experiment, we asked jobseekers their expected search duration and earnings conditional on finding a job and then their expectations for another jobseeker who had better numeracy skills but was otherwise identical to themselves. Jobseekers expect that the other hypothetical jobseeker will search for 0.74 fewer months than themselves (24% of the mean, $p = 0.02$) and earn 118 USD more (13% of the mean, $p < 0.001$). Between-jobseeker estimates of this relationship produce a similar results (Table C.7).

3 Tight Experiment: Effects on Beliefs and Directed Search

The previous section’s conceptual framework and descriptive evidence suggest that jobseekers’ beliefs about their skills might influence their search direction and outcomes. Here we use the tight experiment to study how jobseekers’ skill beliefs and search direction react to new information about their skills. This experiment collects rich data on beliefs and unique measures of skill-directed search using jobseekers’ choices between jobs with different skill demands.

3.1 Experimental Design and Intervention

We ran the experiment during day-long job search workshops. We randomized treatment at the workshop level to avoid spillovers between jobseekers. We assigned 17 workshops to treatment and 17 to control. 373 jobseekers attended the workshops. Our main analysis uses the 278 who have a unique comparative advantage. Treatment assignments are balanced on baseline covariates, for both the 373 and the 278 jobseekers (Table D.1). Jobseekers knew they were participating in a research study but not an experiment, so treatment-control differences in outcomes should not reflect experimenter demand effects. All jobseekers received reimbursement for transport and the time spent at the workshop.

The timeline of the day is shown in Figure D.1. Jobseekers first completed a pre-treatment survey, in which we defined communication and numeracy skills and the concept of quintiles. We asked jobseekers their beliefs about which quintile their general communication and numeracy skills are in, relative to the reference group. They then took assessments of their numeracy, communication, and concept formation skills and completed another survey about their perceived performance on the assessments.

Treated jobseekers then received a report describing the assessments and their performance (Figure 1). For each skill, the report shows the quintile in which the jobseeker ranked on each assessment, compared to other jobseekers in the reference group. Treated

jobseekers watched a video that explained the skill assessments and how to interpret the report, particularly the quintiles. The video encouraged them to think about what jobs will value their skills but did not encourage applying to any specific types of jobs.

Jobseekers in the control group did not receive a report. They watched a control video, which was a strict subset of the treatment video. It contained the parts of the treatment video that explain the assessments, to hold constant any effect of having better information about communication or numeracy skills. It also contained the encouragement to jobseekers to think about what jobs will value their skills, to hold constant the general idea of skill-directed job search. Both groups took the same assessments and answered the same skill-focused survey questions, so any priming effects about the importance of skill levels and skill match are held constant across treatment groups.

To facilitate comprehension, we intensively piloted reports and videos and gave jobseekers time to ask questions during and after the video. After the video, we asked treated jobseekers three understanding checks: 99% and 96% correctly reported the quintiles they scored for respectively numeracy and communication, and 98% correctly reported the skill in which they scored higher. Belief updating does not differ by conceptual formation scores, suggesting general cognitive ability does not limit processing of this information.

The report is designed to provide information only to the jobseekers themselves, not to prospective employers. The report does not include the jobseeker’s name or any identifying information and has no Harambee branding. We show in Section 6 that information acquisition by firms is unlikely to explain the results of the experiment.

Appendix D contains a detailed description of the workshops and links to the videos.

3.2 Specification

We estimate the effects of receiving information about one’s relative skill ranking:

$$Y_{id} = T_d \cdot \beta + \mathbf{X}_{id} \cdot \Gamma + \epsilon_{id}. \quad (3)$$

β , the average treatment effect, is the main object of interest. Y_{id} is the outcome for jobseeker i assessed on date d , T_d is a treatment indicator, and $\mathbf{X}_{id}$ is a vector of prespecified baseline covariates.¹⁵ We use heteroskedasticity-robust standard errors clustered by workshop date, the unit of treatment assignment.

We also test whether treatment effects are larger for jobseekers whose baseline beliefs are misaligned with their comparative advantage, because these jobseekers receive less

¹⁵ $\mathbf{X}_{id}$ contains age; a dummy for being female; dummies for only high school education, having a post-secondary certificate, and for having a post-secondary degree; dummies for each of the skill quintiles for both numeracy and communication skills; a pre-treatment value of the outcome Y_{id} where available; and block fixed effects, to account for the fact that we randomize treatment within blocks of 4 sequential days.

Figure 1: Sample Report

REPORT ON CANDIDATE COMPETENCIES

-Personal Copy-


This report contains results from the assessments you took today. These results can help you learn about some of your strengths and weaknesses and inform your job search.

You completed assessments on English Communication (listening and reading comprehension) and Numeracy today.

1. The Numeracy test measures various maths abilities.
2. The Communication test measures English language ability through listening and reading comprehension.

Your results have been compared to a large group of young South African job seekers who have a matric certificate, have completed school in rural areas or townships around Johannesburg and have completed the same assessments.

You scored in the **FIRST GROUP** of these candidates for Numeracy and the **THIRD GROUP** for Communication.



Note: **Figure 1** shows an example of the reports given to treated jobseekers. Each report contains the jobseeker's assessment results but no identifying information (name, national identity number, etc.) and does not include any Harambee branding. "Completed school in rural areas or townships" is a common proxy in South Africa for attending school in a low-income area.

information from treatment. We estimate models of the form:

$$Y_{id} = T_d \cdot \alpha^{misaligned} + T_d \cdot Aligned_{id} \cdot \alpha^{diff} + Aligned_{id} \cdot \delta + \mathbf{X}_{id} \cdot \Gamma + \epsilon_{id}. \quad (4)$$

$Aligned_{id}$ is an indicator for jobseekers whose pre-treatment beliefs about their comparative advantage on the assessments match their assessment results (measurement details in Appendix C.1). We report the average treatment effect for jobseekers with misaligned baseline comparative advantage beliefs, $\alpha^{misaligned}$; the average treatment effect for jobseekers with aligned baseline comparative advantage beliefs, $\alpha^{misaligned} + \alpha^{diff}$; and the difference between the treatment effects for the two subgroups, α^{diff} .

For both specifications, we restrict our analysis on the 278 individuals with a clear comparative advantage in one of the skills, of whom 139 were treated. This is necessary to cleanly define job choices aligned with comparative advantage in skills, as discussed on page 13. However, including these jobseekers in the sample and classifying all their job search choices as misaligned with their comparative advantage produces qualitatively similar treatment effects on skill beliefs (Table E.1) and on search direction (Table E.2).

As in any heterogeneity analysis, $Aligned_{id}$ may be correlated with other jobseeker-level characteristics, complicating interpretation of α^{diff} . This is a relatively minor concern in this specific analysis because having an aligned comparative advantage belief at baseline is unrelated to gender, age, employment, and work experience. It is related to the communication and numeracy assessment scores. But all these variables jointly explain only 18% of the belief variation (Table C.4). Thus, $Aligned_{id}$ likely captures heterogeneity by baseline comparative advantage beliefs more than other by other characteristics.

Both estimating equations, all baseline covariates, and most outcome measures are prespecified at <https://www.socialscisceregistry.org/trials/10000/>. We describe the relationship between the preanalysis plan and our final analysis in Appendix K.

3.3 Information About Skills Aligns Beliefs with Comparative Advantage

In the control group, 47.5% of jobseekers have aligned comparative advantage beliefs: they believe they rank in a higher quintile for the skill in which they score in a higher quintile on our assessments. Treated jobseekers are on average 13.5pp more likely to report aligned beliefs, a 28% increase on the control group mean (Table 2, column 1, $p < 0.001$). This measure captures their beliefs about their skills in general, not their results on our specific assessments. Appendix C.1 describes in detail how we measure beliefs.

Most of the treatment effect is driven by jobseekers with misaligned comparative advantage beliefs at baseline. Treatment increases the proportion of these jobseekers with aligned comparative advantage beliefs by 21pp (44% of the control mean, $p < 0.001$). This is the estimate of $\alpha^{misaligned}$ from equation (4), shown in Table 2, column 2, row 1.

Table 2: Treatment Effects on Beliefs About Skills - Tight Experiment

	Aligned CA belief		Fraction aligned beliefs	
	(1)	(2)	(3)	(4)
Treatment	0.135*** (0.035)	0.208*** (0.050)	0.078*** (0.026)	0.036 (0.028)
Treatment $\times$ Aligned CA belief (bl)		-0.137 (0.082)		0.080 (0.050)
Aligned CA belief (bl)		0.586*** (0.079)		-0.050 (0.046)
Treatment effect: Aligned CA belief (bl)		0.072 (0.058)		0.116*** (0.041)
Control mean	0.475	0.475	0.183	0.183
Observations	278	278	278	278

Notes: Table 2 shows that treatment aligns jobseekers' beliefs about skills with their assessed skills in the tight experiment. "CA" stands for comparative advantage and "bl" stands for baseline. Columns indicate different outcomes: a dummy indicating if a jobseeker's belief about her CA in skills is aligned with the assessed CA (cols. 1-2), and the fraction of skills where her believed and assessed quintile are equal (cols. 3-4). Cols. 2, and 4 show treatment effect heterogeneity by whether jobseekers had aligned CA beliefs at baseline. Control variables are defined in footnote 15. Standard errors clustered at the treatment-day level shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In contrast, treatment has modest effects on jobseekers with aligned beliefs at baseline. The proportion of this group with aligned beliefs increases by only 7.2pp ($p = 0.214$). This is the estimate of $\alpha^{misaligned} + \alpha^{diff}$ from equation (4), shown in Table 2, column 2, row 4. The difference in treatment effects between jobseekers with aligned and misaligned baseline beliefs is large – 13.7pp, or 28.8% of the control mean – but not quite statistically significant ($p = 0.105$). This is the estimate of α^{diff} , shown in Table 2, column 2, row 2.

Treatment shifts jobseekers' beliefs about their skill levels as well as their beliefs about skill comparative advantage. In the control group, 18% of jobseekers' beliefs about their skill levels match their assessed skill quintiles. Treatment increases this by 7.8pp, a 43% increase (Table 2, column 3, $p = 0.005$). We show heterogeneous effects by baseline comparative advantage beliefs (column 4) but do not focus on them, because baseline comparative advantage beliefs do not reflect individuals' scope to learn about their skill levels.

We find similar patterns using more continuous measures of (mis)alignment between assessed skills and beliefs about skills, both for the outcome variables and for the baseline alignment measure used to estimate heterogeneous treatment effects (Table E.4). Updating of comparative advantage beliefs is at least partially explained by learning about the distribution of skills in the reference population (Appendix C.2). Treatment updates un-

derconfident beliefs more than overconfident beliefs (Table C.8), matching findings in other research (e.g. Eil & Rao 2011). Belief updating does not differ by gender (Table J.3) so we pool genders for the rest of the experimental analysis.

3.4 Job Search with Better-Aligned Beliefs about Skills

We show here that that treatment shifts jobseekers' search toward jobs that align with their assessed comparative advantage across four measures of skill-directed search.


Job choice task: We design a novel incentive-compatible job search task in which participants make 11 choices between pairs of job advertisements/vacancy postings. In each pair, one job had been coded by recruiters as requiring more numeracy skills and one as requiring more communication skills. Participants were shown each pair of advertisements, given time to read them, and asked to select one to apply to. Figure 2 shows an example pair and Appendix Table D.2 shows the full list of job titles. Participants viewed the 11 job advertisement pairs in random order.

The postings in the task are based on real job advertisements on SAYouth.mobi. All jobseekers in the study used the platform, so choices between vacancies represented real-life choices they often made. To select the specific postings, we considered entry-level vacancies in Johannesburg with no specialized education requirements posted on SAYouth.mobi. Among these, we selected a longlist of 28 vacancy postings with a clear numeracy or communication skill requirement, recognizing that not all jobs on the platform have such requirements. We asked 13 recruitment professionals with experience hiring for entry-level roles to rate each vacancy on required communication and numeracy skills, transparency of skill requirements, expected wage, and overall desirability. We then created 11 pairs of jobs. Within each pair, both jobs had similar expected wages and desirability but one job required more communication skill and one required more numeracy skill. Averaging across all pairs, recruiters scored the job we defined as numeracy-heavy as needing 2.7 standard deviations higher numeracy skills but similar on other aspects. For example, the average within-pair difference in expected wages was $< 5\%$ of the mean wage. We removed the employer name and location and standardized length and format.

The job pairs thus differed only in their skill demands, while treatment gave information only on jobseekers' relative ranks on communication and numeracy skills. This design allows us to measure how jobseekers' skill-directed job search changes in response to information about their skills, holding other job attributes constant. In non-experimental datasets, jobs with different skill demands may also differ on unobserved characteristics that drive observed search direction, making it difficult to study this question.

We incentivized jobseekers to respond truthfully in the job choice task in two ways.

Figure 2: Sample Pair of Jobs from Job Choice Task



First, one pair contained live advertisements for jobs at a partner firm. We told jobseekers we would submit their application to the job they chose from this pair but did not tell them which pair it was. Second, we told jobseekers that after the workshop we would send them recommendations for entry-level job titles matching their choices in the task.

Our main outcome is the share of the 11 pairs in which the jobseeker chose the job that required the skill aligned with her measured comparative advantage. In the control group, this share was 55%. Treatment increases this share by 3.7pp (Table 3, column 1, $p = 0.285$). This average effect hides important heterogeneity. For jobseekers with misaligned baseline beliefs about their comparative advantage, treatment increases the share of aligned choices by 8.8pp, 16% of the control mean (column 2, $p = 0.028$). Treatment does not change search direction for jobseekers with baseline aligned beliefs.

We see a consistent pattern of heterogeneity for comparative advantage beliefs and search direction: treatment effects on both outcomes are driven by jobseekers with unaligned baseline beliefs about their skill comparative advantage. This pattern is also consistent with the predictions from our conceptual framework.

The search direction results suggest that jobseekers can interpret information in postings about the relative skill demands of different jobs. However, they might not perfectly interpret jobs' skill demands. We tested this by explicitly revealing relative skill demand for the last two pairs of jobs (recall that the order in which jobseekers see the 11 job pairs was randomized). Jobseekers with initially misaligned beliefs align their search for job choices both with and without revealed skill requirements, but do so substantially more

when skill requirements are revealed (Table E.8). This shows they have limited information about job skill requirements but enough for some skill-directed job search.

Application data from online search platform: We linked each jobseeker to their profile on SAYouth.mobi and observed their on-platform job search. In the 30 days after the workshop, the average participant in both the treatment and control groups started 15 job applications on the platform, slightly higher than in the preceding 30 days (Table H.2, column 5). The platform does not consistently record if applications are completed.

We classified vacancies by skill demand where possible. There were 69,000 vacancies posted on SAYouth.mobi in Johannesburg during our study period. This captures all online job postings for Johannesburg as SAYouth.mobi automatically scrapes data from all other online job boards. Of these vacancies, we classified 14% as communication-heavy and 13% as numeracy-heavy jobs, suggesting roughly equal demand for both skills among jobs posted online.¹⁶ Post-treatment, 23% of applications started by jobseekers in our sample were to jobs classified as requiring communication or numeracy skills.

Treated jobseekers were more likely to start applications to jobs demanding skills that matched their measured comparative advantage. To show this, we calculate each jobseeker's number of applications to vacancies coded as requiring the skill aligned with their comparative advantage, subtract the number of applications to vacancies coded as requiring the opposite skill, and divide this difference by the number of applications to vacancies coded as requiring either skill. Treated jobseekers start 6.3pp more aligned applications than non-aligned applications (Table 3, column 3, $p = 0.010$). Again, this average effect is driven by jobseekers with misaligned baseline comparative advantage beliefs, who have an effect of 8.9pp (column 4, $p = 0.074$). We obtained these data directly from the platform after the experiment, so experimenter demand effects are unlikely.

Clicks on links to real jobs: Our third measure captures whether jobseekers can conduct skill-directed job search using only job titles. We sent jobseekers three text messages with links to real job opportunities on SAYouth.mobi about a week after the workshop. We sent, in random order, one numeracy job, one communication job, and one job aligned with the skill demand of the majority of their choices in the job-choice task. The messages contained a greeting linking the message to the workshop they attended, a note that we found a job opportunity of interest within commuting distance of the workshop venue, the job title, the link to apply, and the #SAYouth hashtag (which Harambee regularly uses

¹⁶We classify job adverts that contain one of 20 numeracy- or 20 communication-heavy job titles in the advert title or job description as requiring the respective skill. The lists of skill-specific job titles was created by the recruitment staff who classify vacancies for the job choice task. Many job postings omit wage information so we cannot construct wage-based measures of demand for communication and numeracy skills in this market.

Table 3: Treatment Effects on Search Direction - Tight Experiment

	% aligned (job choice)		Δ % aligned platform apps		Δ SMS click rate		Δ planned apps (w)		Aligned search index	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment	0.037 (0.034)	0.088** (0.038)	0.063** (0.023)	0.089* (0.048)	0.071 (0.061)	0.157 (0.096)	1.420 (1.261)	4.746** (1.763)	0.269** (0.103)	0.603*** (0.140)
Treatment $\times$ Aligned CA belief (bl)		-0.104** (0.039)		-0.048 (0.080)		-0.163 (0.128)		-6.629** (2.651)		-0.648*** (0.205)
Aligned CA belief (bl)		0.165*** (0.035)		0.013 (0.049)		0.100 (0.102)		8.697*** (2.412)		0.727*** (0.153)
Treatment effect: Aligned CA belief (bl)		-0.016 (0.034)		0.041 (0.044)		-0.005 (0.081)		-1.883 (1.807)		-0.045 (0.131)
Control mean	0.550	0.550	0.007	0.007	-0.032	-0.032	4.331	4.331	-0.000	-0.000
Observations	278	278	278	278	278	278	278	278	278	278

Notes: Table 3 shows that informing jobseekers about their relative comparative advantage in skills aligns their search direction with their assessed comparative advantage in the tight experiment. “CA” stands for comparative advantage and “bl” stands for baseline. Aligned job search is defined as directing search effort toward jobs that mostly require the skill that aligns with jobseekers’ assessed CA. Columns indicate different outcomes: the percentage of 11 incentivized job choices that are aligned with the measured CA of the jobseeker (cols. 1–2), the difference between the percentage of aligned and non-aligned applications on the online job search platform SAYouth.mobi (cols. 3–4), the difference in link click rates between aligned and non-aligned jobs sent to job seekers via text message (cols. 5–6), the difference between aligned and non-aligned planned applications for the 30 days after the workshop (cols. 7–8), and an inverse-covariance weighted average of the search alignment measures displayed in cols. 1–8 following [Anderson \(2008\)](#) (cols. 9–10). Even columns show heterogeneity by whether individuals have aligned CA beliefs at baseline. Control variables are defined in footnote 15. All winsorized variables (w) are winsorized at the 99th percentile. Standard errors clustered at the treatment-day level shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

in messages to platform users). We tracked whether jobseekers clicked on these links.

We again find treatment increases skill-directed job search. Treatment increases the difference in click rates between jobs that are aligned and misaligned with the jobseeker’s comparative advantage by 7pp (Table 3, column 5, $p = 0.254$). The effect is again driven by jobseekers with initially misaligned beliefs. However, estimates are relatively imprecise because this measure uses 3 choices per jobseeker, compared to 11 in the job choice task.

Planned applications to numeracy and communication jobs: After the treatment, before the job choice task, we surveyed participants about the number of applications they planned to send to communication-heavy and numeracy-heavy jobs in the next 30 days. We calculate the number of planned applications to jobs aligned with their assessed comparative advantage, minus the number of planned applications to jobs focused on the other skill. Treatment increases this by 1.42 applications, a 33% increase on the control mean (Table 3, column 7, $p = 0.269$). This average effect is driven by a 4.8 application effect for jobseekers with misaligned baseline comparative advantage beliefs (column 8, $p = 0.011$). The result reflects both a slight increase in planned aligned applications and a reduction in planned misaligned applications. This measure is more susceptible to experimenter demand effects but produces similar results to the other measures.

Search direction index: We combine these four measures of skill-directed search into an index to avoid multiple hypothesis testing and to increase power (Anderson, 2008). We see a large, positive treatment effect of 0.27 standard deviations (Table 3, column 9, $p = 0.014$). This effect is entirely driven by jobseekers with misaligned baseline comparative advantage beliefs, for whom we observe a shift in search direction of 0.6 standard deviations (Table 3, column 10, $p < 0.001$). For this index, treatment closes roughly 80% of the gap between jobseekers with aligned versus misaligned baseline beliefs: treated jobseekers with misaligned baseline comparative advantage beliefs have an average skill-directed job search index of 0.60 (column 10, row 1), while control jobseekers with aligned baseline comparative advantage beliefs have an average of 0.73 (column 10, row 3).

3.5 Additional Results and Robustness Tests

Beliefs about returns to directed search: The model predicts that new information about jobseekers’ relative skill ranks will also change their expectations about the relative returns to searching for jobs that require different skills. To test this, we measure jobseekers’ expectations about the outcomes of applying to communication- and numeracy-heavy jobs, both in the job choice task and in their planned search after the workshop. For example, for a jobseeker with comparative advantage in numeracy, we construct her expected

wage for numeracy-heavy jobs minus her expected wage for communication-heavy jobs.

Treatment increases expected returns to most measures of skill-directed search (Table F.1), although some results are imprecisely estimated. Jobseekers' expectations about job-specific wages also predict choices in the job choice task (Table F.3). These results are consistent with the model. Appendix F gives full details of the measurement and results.

Robustness checks for beliefs and search results: Treatment effects estimates are qualitatively similar when we include jobseekers without unique comparative advantage in the sample (Tables E.1 and E.2). Hypothesis test results are robust to using a wild cluster bootstrap to account for having only 34 clusters and to accounting for multiple hypothesis testing following Benjamini et al. (2006) (Table E.3). Treatment effect estimates are robust to using more continuous measures of skill belief alignment (Table E.4) and to controlling for baseline 'confidence' – the difference between believed and assessed skill quintiles – and its interaction with treatment (Table E.7).

4 Big Experiment: Effects on Beliefs, Search & Labor Market Outcomes

Having shown that information about skills can shift comparative advantage beliefs and search direction, we now test if this information can also improve individual labor market outcomes in an experiment with 4,389 participants and a follow-up period of 3.5 months.

4.1 Sample

The big experiment took place in the same location in 2016/17, five years before the tight experiment. Recruitment was from the same population: active jobseekers in the database of our partner Harambee, described in Section 2.3.

Table A.2 shows both samples have similar labor market participation. In the tight/big experiment, 33/37% had done some work or income-generating activity in the past seven days and 96/97% were actively searching for work. The average jobseeker submitted 10/9 job applications and spent 14/17 hours searching for work in the last seven days.

The samples differ in some demographic characteristics. Jobseekers in the big experiment, relative to the tight experiment, were slightly more likely to be male (38 vs 33%) and younger (24 vs 26), reflecting changes in Johannesburg's demographics (SSA, 2016, 2022). A similar proportion had finished high school, but slightly fewer in the big experiment had university degrees (17 vs 22%) and slightly more had other, shorter tertiary qualifications (22 vs 15%), consistent with national trends. Only 9% had ever held a permanent or long-term job (vs 25% in the tight experiment), perhaps because they were younger.

Our key results are robust to accounting for these differences. Table E.3 shows that the main treatment effects in both experiments are robust to reweighting the two samples to have the same distribution of baseline demographics, education, and employment.

4.2 Experimental Design

We assigned 2,114 jobseekers in 27 workshops to the treatment group and 2,274 jobseekers in a different 27 workshops to the control group. To avoid spillovers, treatment was administered at the workshop (day) level. Treatment groups are balanced on covariates for both the full sample and the 96% recontacted for the endline survey (Table D.3). Endline response rates are balanced across treatment groups (Table D.4).

The big and tight experiments have deliberately similar designs. The big experiment was also run during Harambee workshops that included skill assessments, surveys, and basic job search advice. Jobseekers in treated workshops received a report about their skill assessment results (Figure D.3); jobseekers in control workshops did not. Participants knew they were participating in a study but not that treatment differed by day, so treatment-control differences in outcomes should not reflect experimenter demand effects. Appendix D provides detailed descriptions of the workshops and interventions including figures showing the sequence of events in the workshops (Figures D.1 & D.2).

The experiments differ in four ways. First, the big experiment was designed to measure effects on labor market outcomes, so we studied a larger sample and collected post-treatment data later: on average 3.5 months after treatment.¹⁷

Second, in the big experiment we assessed and gave treated jobseekers information on six skills, rather than two. This is more similar to real-world settings, where jobseekers direct search across a broader range of job types with different skill demands. In settings like schools or job centers, they often receive and must process information about their relative rank on multiple skills. In contrast, the tight experiment used fewer skills to allow cleaner definitions of comparative advantage and clearer tests for skill-directed search.

We assessed communication, numeracy, concept formation, focus, grit, and planning. (Details in Appendix B). The reports given to treated jobseekers showed results for all six assessments and information on what traits they measure (Figure D.3).¹⁸ The six assessments differentiate jobseekers horizontally more than vertically because assessment results are weakly correlated across skills within candidate and do not allow an easy aggregate ranking of candidates. 13 of the 15 pairwise correlations are < 0.3 (Table B.3). Most jobseekers' reports showed substantial variation across skills: 85% had at least one top tercile but only 1.7% had all six top terciles and 58% had both top and bottom terciles.

Third, we reported assessment results in terciles, not quintiles. The coarseness of ter-

¹⁷We used phone surveys lasting on average 25 minutes, compensating respondents with mobile phone airtime. Phone and in-person surveys in this setting deliver similar labor market data (Garlick et al., 2020).

¹⁸In piloting, jobseekers easily recognized the traits linked to the soft skills measures. For example, they did not use the term "grit" but did refer to the persistence required for repetitive, boring jobs.

ciles relative to quintiles and using six rather than two skills on the reports means that only 23% of jobseekers have a unique skill comparative advantage, compared to 75% in the tight experiment. In the tight experiment, the skill-directed search measures could not be sensibly defined for jobseekers without a unique skill comparative advantage, so we dropped them from the main analysis and included them only in robustness checks. In the big experiment, we keep these jobseekers in all analyses because the skill-directed job search measures in this experiment can be defined for them.

Fourth, there are logistical differences that are unlikely to affect the core mechanism activated by the treatment. In the big experiment, each workshop had more candidates. The briefings after candidates receive the report were delivered in person instead of by video, but with a script to ensure consistency across sessions. The job search assistance – advice writing CVs and cover letters – was at the beginning of the day in the big experiment and the end of the day in the tight experiment.

4.3 Specification

We estimate average treatment effects using the following specification:¹⁹

$$Y_{id} = T_d \cdot \beta + \mathbf{X}_{id} \cdot \Gamma + \varepsilon_i \quad (5)$$

Y_{id} is the outcome of interest for jobseeker i assessed on date d and $\mathbf{X}_{id}$ is a vector of pre-specified control variables.²⁰ The object of interest is the average treatment effect β . As in the tight experiment, we cluster standard errors at the level of treatment (assessment days) and we combine families of outcomes in inverse covariance-weighted average indices to limit the number of tested hypotheses, following [Anderson \(2008\)](#). Our main results are robust to using a wild cluster bootstrap and to adjusting for multiple hypothesis testing, following [Benjamini et al. \(2006\)](#) (Table E.3).

4.4 Information About Skills Aligns Beliefs with Assessed Comparative Advantage

Receiving information about skills substantially increases alignment between measured skills and skill beliefs. To show this, we ask candidates in which tercile they believe they ranked for each of the communication, numeracy, and concept formation assessments (see Appendix C.1 for measurement details). We define a jobseeker as having an aligned comparative advantage belief if she believes she scored highest in the skill(s) in which she

¹⁹Most of the outcome variable definitions, inference methods, and covariates were prespecified, as for [Carranza et al. \(2022\)](#). However, comparative advantage beliefs and aligned search were not prespecified as outcomes. Our analysis of these outcomes should be viewed as ex post in the big experiment but pre-specified in the tight experiment. See Appendix K for details.

²⁰ $\mathbf{X}_{id}$ contains baseline assessment results, self-reported skills, education, age, gender, employment, earnings, job offers, time and risk preferences, self-esteem, baseline values for the outcome where available, and fixed effects for the blocks of days within which treatment was randomized.

Table 4: Treatment Effects on Skill Beliefs and Search Direction - Big Experiment

	Beliefs		Search direction
	Aligned CA belief (1)	Fraction aligned beliefs (2)	Aligned search (3)
Treatment	0.139*** (0.011)	0.142*** (0.009)	0.050*** (0.010)
Control mean	0.196	0.388	0.165
Observations	4118	4195	4205

Notes: Table 4 shows that informing jobseekers about their comparative advantage in skills aligns their skill beliefs with their assessed skills and shifts self-reported search direction in the big experiment. “CA” stands for comparative advantage. Columns indicate different outcome variables: dummy for aligned CA belief (col. 1), fraction of aligned skill tercile beliefs across communication, numeracy, and concept formation (col. 2), and a dummy indicating if self-reported search direction is aligned with assessed skill CA (col. 3). Control variables are described in footnote 20. Standard errors clustered at the treatment-day level shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

actually scored highest. If she has tied top terciles, we require she also believes she has tied top skill terciles.

Only 20% of control group jobseekers have aligned comparative advantage beliefs using this definition. Treatment increases this by 13.9pp, a 70% increase (Table 4, column 1, $p < 0.001$). This is in line with the effects on beliefs observed straight after treatment in the tight experiment.²¹ This shows that jobseekers’ updated beliefs persist over 3.5 months and that jobseekers are able to process more complex information about assessment results in multiple skills.

Treatment also increases the fraction of beliefs about skill levels that align with assessment results by 14.2pp, a 37% increase from the control mean of 38.8pp (Table 4, column 2, $p < 0.001$). As in the tight experiment, we observe no gender heterogeneity in baseline beliefs or belief updating, so we report all analysis pooling genders (Table J.4).

4.5 Job Search with Better-Aligned Comparative Advantage Beliefs

We collect a self-reported measure of alignment between jobseekers’ search direction and skill comparative advantage. We ask jobseekers to think about the types of jobs for which they are applying and what skills these jobs demand and then to rank the following skills from most to least important: numeracy, communication, problem-solving, and soft skills.

²¹The effects on aligned comparative advantage beliefs are almost identical across the tight and big experiments. But the control group mean is higher in the tight than big experiment because the big experiment uses three skills, creating more ways for believed and assessed comparative advantage to differ. In both experiments, the control group’s alignment rate is only slightly higher than random guessing, accounting for the possibility of ties: 20 versus 15% in the big experiment and 48 versus 40% in the tight experiment.

Aligned search is a dummy indicating that a jobseekers' highest ranked skill on this measure matches the skill in which she has an assessed comparative advantage.²² Treatment increases this measure from 17% to 22% (Table 4, column 3, $p < 0.001$). This result on self-reported skill-directed job search, measured 3.5 months after treatment, is qualitatively consistent with the four measures of skill-directed search in the tight experiment, collected in the first month after treatment.²³

4.6 Improved Labor Market Outcomes

The results on beliefs and search from the two experiments provide strong evidence that giving jobseekers information about their skill assessment results better aligns their search with their assessed skill comparative advantage. We now test the model prediction that the change in search direction should improve jobseekers' labor market outcomes.

Treatment has limited effects on employment quantity: Treatment effects on the number of job offers and probabilities of employment in the first and second months after the workshop and the week before the endline survey are positive but mostly small and not statistically significant (Table 5, panel A, columns 1–4). Our prespecified index of employment quantity increases by 0.05 standard deviations (column 5, $p = 0.166$).

Treatment substantially increases employment quality: Treatment increases earnings in the seven days before the endline by 6.52 USD (Table 5, panel B, column 1, $p = 0.020$). This is equivalent to a 26% increase from the control group mean or to moving from the 75th to 79th percentile of the unconditional earnings distribution.

We code earnings as zero for all non-employed jobseekers, so this variable has a mass point at zero and a continuous distribution above zero. Treatment effects on such measures can be sensitive to scaling choices (Mullahy & Norton, 2022) so we report several robustness checks. First, effects are weakly positive on the entire earnings distribution (Figure E.1). Second, treatment raises earnings by 9.4 USD without winsorization, 0.1 standard deviations, 12 log points or 13 inverse hyperbolic sine points, demonstrating robustness across a range of scalings (Table E.5, columns 1 and 3-5). Third, restricting the sample to the employed produces qualitatively similar treatment effects relative to the control mean (Table E.5, columns 6-7).²⁴ The earnings effects are within the range of

²²To construct this dummy, we remove soft skills from the jobseeker's ranking and equate problem solving with the concept formation assessment. For individuals with a tie for highest skill, our main aligned search measure is set to zero. However, results are robust to recoding this measure to one when the jobseeker said that any one of her tied highest skills is most important for the jobs to which she was applying.

²³The big experiment effect converts to a 0.14 standard deviation increase in self-reported skill-directed search, compared to a 0.27 standard deviation effect on the skill-directed search index in the tight experiment (Table 3, column 1). The effect size on the index is larger because the index averages over four measures, producing a smaller standard deviation and hence larger standardized effect size.

²⁴Restricting analysis to the employed can create a sample selection problem when treatment affects the

Table 5: Treatment Effects on Labor Market Outcomes - Big Experiment

Panel A: Work quantity					
	Job offers (w)	Worked m1	Worked m2	Worked last 7d	Index
	(1)	(2)	(3)	(4)	(5)
Treatment	0.015 (0.017)	0.023* (0.013)	0.007 (0.015)	0.009 (0.013)	0.045 (0.032)
Control mean	0.182	0.465	0.437	0.309	-0.000
Observations	4140	4201	4204	4204	4205

Panel B: Work quality				
	Earnings (w)	Hourly wage (w)	Written contract	Index
	(1)	(2)	(3)	(4)
Treatment	6.517** (2.712)	0.295** (0.130)	0.017* (0.010)	0.085** (0.035)
Control mean	25.424	1.267	0.120	0.000
Observations	4196	4183	4184	4206

Notes: Table 5 shows that the treatment improves employment quality but not necessarily employment quantity in the big experiment. Panel A shows effects on employment quantity. The columns indicate different outcomes: the winsorized number of job offers in the last 30 days (col. 1), a dummy indicating any work for pay in month 1 after treatment (col. 2), a dummy indicating any work for pay in month 2 after treatment (col. 3), a dummy indicating any work for pay in the seven days before the endline survey (col. 4), and an inverse covariance-weighted average of the four employment quantity measures (col. 5). Panel B shows effects on employment quality. The columns indicate different outcomes: winsorized earnings in the last seven days (col. 1), winsorized hourly wages in the last seven days (col. 2), a dummy indicating a written contract (col. 3), and an inverse covariance-weighted average of the three employment quality measures (col. 4). Employment quality measures are coded as zeroes for the non-employed. Control variables are described in footnote 20. All monetary figures are reported in 2021 USD in purchasing power parity terms. All winsorized variables are winsorized at the 99th percentile. Standard errors clustered at the treatment-day level shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

effect sizes reported in related research, although toward the top of that range.²⁵

Treatment increases hourly wages by 0.30 USD or 23% of the control mean of 1.27 USD (Table 5, panel B, column 2, $p = 0.028$). Treatment also increases the probability of having a written contract by 1.7 percentage points from the 12% mean in the control group (Table

probability of employment. But, in this case, treatment has little impact on employment, so it is unsurprising that the results are qualitatively robust to different ways of handling the earnings of the nonemployed. Treatment effects for the log and inverse hyperbolic sine specifications in the sample of employed jobseekers are slightly larger, potentially due to these specifications' scale-dependence (Mullahy & Norton, 2022).

²⁵For example, interventions that combine learning about skills with the ability to signal skills increase earnings by 11-34% (Abebe et al., 2021b; Bassi & Nansamba, 2022; Carranza et al., 2022). Bandiera et al. (2023) find an 11% change in earnings from a matching intervention that they attribute to effects on beliefs about labor market prospects. Using a more model-based approach, Guvenen et al. (2020) estimate that moving from the bottom to the top decile of skill match quality increases earnings by 11 percent.

5, panel B, column 3, $p = 0.095$). Written contracts define formal jobs in this context and are valued by workers. Our prespecified index of employment quality increases by 0.085 standard deviations (Table 5, panel B, column 4, $p = 0.019$).

Heterogeneous treatment effects: We report heterogeneous treatment effects for jobseekers with aligned and misaligned baseline comparative advantage beliefs, as in the tight experiment. However, we interpret these effects with caution because this is a noisier proxy for the “dose” of information delivered by the big than tight experiment: we observe beliefs about only three of the six skills in the big experiment and the additional skills allow many more ways for assessed and believed comparative advantage to differ.

Subject to these caveats, results of this experiment are roughly consistent with larger effects for jobseekers who get more information from treatment. Treatment effects on labor market outcomes are driven by jobseekers with misaligned baseline comparative advantage beliefs (Table E.10). Although none of the differences is statistically significant at conventional levels, some are large. For example, the effects on weekly earnings are 7.5 and 2.8 USD for jobseekers with respectively misaligned and aligned baseline comparative advantage beliefs (column 6). Similarly, the treatment effect on skill-directed job search is driven by jobseekers with misaligned baseline beliefs (Table E.6, column 3), although this does not hold for the effect on comparative advantage beliefs (column 1).

4.7 How Does Skill-Directed Job Search Improve Labor Market Outcomes?

Here we explore possible economic mechanisms that could link skill-directed search to better-paid and higher-quality jobs but not to a higher employment probability. Given the nature of our data, this subsection is more speculative than the rest of the paper.

We begin by motivating our focus on search-based mechanisms. Treatment effects on earnings and formality are driven by wage jobs that started after treatment, rather than self-employment or wage jobs that started before treatment (Table E.9). This suggests that self-employment and bargaining with existing employers are less important mechanisms than search.²⁶ Given this pattern, we evaluate two non-mutually exclusive mechanisms centered on two standard concepts in the job search literature: **job offer probabilities** and **wage offer distributions**. Returning to our model from Section 2.1, we assume there are multiple job types indexed by j and searching for each job type has an associated wage offer distribution $F_j(w)$ and job offer probability P_j . Each jobseeker’s direction of search effort across different job types generates a wage offer distribution $F(w)$ and job offer probability P that are weighted averages across job types. As in standard job search models, jobseekers accept offers with wages above their reservation wages.

²⁶We don’t distinguish between on- and off-the-job search in this discussion because there are limited differences employed and unemployed participants’ search behavior, in both control and treatment groups.

In the job offer probability mechanism, skill-directed search raises P , so some jobseekers receive multiple offers and can choose the higher-wage offers. Our data are not consistent with this mechanism: treatment effects on the numbers of offers received and refused are both $< 3\%$ of the control group mean ($p > 0.4$) and only 4% of jobseekers receive multiple offers, limiting scope for choice.

In the wage offer distribution mechanism, skill-directed search shifts at least part of $F(w)$ rightward, leading to higher wage offers and raising average earnings for the employed. We do not observe the terms of job offers, so we cannot directly test this. Quantile treatment effects on earnings are consistent with a rightward shift in F (Figure E.1).²⁷

Under the wage offer mechanism, skill-directed job search might generate higher wage offers by raising **job quality** or **match quality**. Higher job quality would occur if some job types pay more to all workers conditional on their skill match, and treated jobseekers applied more to them. Higher match quality would occur if treated and control jobseekers apply to the same job types on average but, within this pool of jobs, treated jobseekers applied more to jobs that better match their skills and hence pay them more.

We find slightly more evidence consistent with higher match quality than with higher job quality. To show this, we construct a list of all job titles that receive an application from any jobseeker in the tight experiment on the [SAYouth.mobi](#) platform. 96.7% of these job titles received applications from both treated and control jobseekers. This suggests a high overlap in the types of jobs to which the groups apply, limiting the scope for a job quality mechanism. Within this overlapping set of job titles, we have already shown that treated jobseekers are more likely to apply for jobs that match their skills. Furthermore, the average jobseeker in the tight experiment sends only 26.7% of applications to jobs that prespecify a fixed wage. This suggests some scope for wage offers to vary with skill match, making a match quality mechanism possible.

We conclude that treatment may raise earnings through more skill-directed search, generating higher wage offers, due to higher match quality. However, our evidence for higher wage offers and match quality is relatively indirect, so we view these as more tentative results than the strong evidence for skill-directed search and higher earnings.²⁸

²⁷To match the null effect on employment, this mechanism also requires either that reservation wages rise, so some of the higher wage offers are rejected, or that the shift in F occurs only above reservation wages. Treatment increases reservation wages in both experiments but the effects are not statistically significant. This may reflect widespread challenges in measuring reservation wages (Feld et al., 2022).

²⁸In principle, part of the treatment effect on labor market outcomes might affect labor market outcomes through come from a feedback loop between beliefs and labor market experience: treatment $\rightarrow$ beliefs about comparative advantage $\rightarrow$ job search $\rightarrow$ labor market outcomes $\rightarrow$ beliefs about comparative advantage. This possibility does not undermine the skill-directed search explanation. It would simply mean that the treatment effects on labor market outcomes reflect both a direct effect and an indirect effect through this feedback loop. The results of the tight experiment show that treatment immediately affects beliefs, before

5 Search Effort and Beliefs about Skill Levels

Can better labor market outcomes be explained by jobseekers being on average overconfident, learning they have lower skills than they thought, and increasing search effort?

We find few treatment effects on search effort in either experiment, suggesting it does not explain the relationship between treatment and labor market outcomes.²⁹ In the tight experiment, treatment effects are close to zero and not statistically significant on six different measures of search effort: a survey question on post-workshop planned applications, time spent on a job search task during the workshop, and four measures of job search on the SAYouth.mobi platform after the workshop (Tables H.2 and H.3). In the big experiment, we estimate precise near-zero effects on self-reported number of applications, time spent and money spent on search (Table H.4).

Treatment does lower jobseekers' beliefs about their skill levels in both experiments (Table H.1, columns 3 and 7). This occurs because treatment shifts jobseekers' beliefs about their skill levels toward their assessed skill levels, and more jobseekers have baseline skill beliefs above their assessed skills than below (columns 4 and 8).

A generalized version of our conceptual framework shows how treatment can lower average beliefs about skills without increasing search effort: because beliefs about skill levels have a theoretically ambiguous effect on search effort. In this framework, jobseekers endogenously choose their total search effort level based on their expected search outcomes, which fall when treatment lowers their believed skill level. The framework predicts two responses to this effect on skill beliefs. There is a substitution effect: jobseekers search less because the expected return to each unit of search effort is lower. There is also an "income" effect: jobseekers search more because more search is needed to achieve the same labor market outcome. The net effect on search effort can be negative, zero, or positive. We show the full conceptual framework and treatment effects in Appendix H.

6 Additional Mechanisms

In this section, we evaluate three other mechanisms that might account for treatment effects on labor market outcomes. These are not mutually exclusive with the directed search mechanism. We find little evidence for any of these three mechanisms.³⁰

Self-esteem: Treatment has near-zero effects on self-esteem in the big experiment, there is time for any feedback loop from effects on labor market outcomes.

²⁹We merely argue shifts in search effort are unlikely to explain the treatment effects on labor market outcomes in this study. We recognize search effort can play an important role in labor market outcomes in other ways, covered in the review by [Mueller & Spinnewijn \(2023\)](#).

³⁰We prespecified testing for human capital investment in both experiments and for all three mechanisms in the big experiment.

in both a text message survey 2–3 days after treatment and the endline phone survey 3.5 months after treatment (Table I.1, columns 1-4), using questions from the Rosenberg (1965) scale. This suggests general beliefs about self-worth do not respond to new skill information and are unlikely to affect search behavior or labor market outcomes.

Human capital investment: We find that jobseekers might make skill investments in response to new information about skills, but that this behavior is unlikely to explain the labor market effects in the big experiment. In the big experiment, treatment has near-zero effects on enrollments in both formal and vocational education, limiting scope for education investment to drive the labor market effects. (Table I.1, columns 5-7).

We find suggestive evidence in the tight experiment that jobseekers might be willing to invest in skills. Treatment reduces willingness-to-pay (WTP) for a numeracy workbook, mostly for people who learn that they have a comparative advantage in numeracy, suggesting that jobseekers might prefer to invest in skills where they are relatively weak. WTP for a communication workbook is unaffected by treatment (Table I.2). Appendix D.1 gives details on the WTP measurement.

Skill information transmission to firms: If jobseekers share assessment results with firms during job applications, this might lead to firm-side learning about jobseekers' skills. We view this mechanism as unlikely to explain labor market outcomes. Labor market effects are driven by the jobseekers who say they did *not* use the report in applications (Table I.3), although we interpret this result cautiously because this analysis conditions on report use, a post-treatment outcome. We have conflicting evidence on how often jobseekers share results with firms: 29% of treated jobseekers in the big experiment report that they ever included a copy of their assessment results with any application in the 3.5 months after treatment but only 0.8% of treated jobseekers receiving this intervention actually included their assessment results in applications to vacancies we created.

The assessment results jobseekers receive are deliberately designed not to be credible to firms. They do not show the jobseeker's name or national identity number, so firms cannot verify that the report is linked to that job applicant. They include no information about Harambee, the source of the assessments, or the value of skills. They are printed in black and white on low-quality paper. None of the 15 hiring managers we interviewed during piloting said they would view these reports as credible.³¹

7 Scale and General Equilibrium Considerations

Job search interventions can generate different results at larger scales due to interactions between jobseekers or between jobseekers and firms (Altmann et al., 2022). Here, we

³¹In contrast, Carranza et al. (2022) study effects of using certificates of assessment results designed to be credible to firms. These are branded by Harambee and contain the jobseeker's name and ID number.

present a brief framework for thinking about scale and general equilibrium implications of the economic mechanisms we study. We provide empirical evidence about parts of this framework, which suggests that the effects of our intervention on jobseekers' labor market outcomes need not be smaller at larger scales. But our experiments are not designed to study general equilibrium. So this section is deliberately brief and speculative.

There are two obvious channels for general equilibrium effects of our intervention. First, jobseekers applying to the same jobs can generate **search congestion**. This occurs when a treatment-induced shift in search behavior increases competition for some jobs or job types, leading directly to lower offer probabilities and potentially to lower wage offers as firms observe higher application volumes. This is unlikely for the type of information intervention we study. Information about skill comparative advantage is inherently differentiated and shifts different jobseekers to apply for different job types, rather than shifting total applications toward one job type. We observe exactly this pattern in the tight experiment. For example, in the job choice task, treatment spreads applications marginally more evenly between communication and numeracy jobs (Table I.4). And treatment does not affect the total numbers of on-platform applications to communication or numeracy jobs. Less differentiated interventions may be more prone to search congestion effects, such as those providing non-differentiated information about high-demand sectors or locations, encouraging search effort, or providing job placement services (e.g. Crépon et al. 2013; Johnston & Mas 2018; LaLive et al. 2022).³²

Second, more skill-directed job search might **raise match quality** or **lower screening costs** for firms, as they receive better-matched applications. These mechanisms can increase aggregate labor demand according to both classic matching models (e.g. Mortenson & Pissarides 1994) and recent experiments (e.g. Algan et al. 2022). We show some suggestive evidence in Section 4 that our intervention raises match quality.

We tentatively conclude that there is no clear evidence that the type of intervention we study would generate smaller effects at larger scales. However, a more conclusive answer would require research designed to evaluate general equilibrium effects.

General equilibrium effects are also unlikely to affect our experiments' internal validity. The big experiment involved only 4,400 people in a city of roughly 8 million people and 2 million employed workers (SSA, 2016); they were sampled from all around the city, not one small geographic area; workshops were spread over seven months; and participating jobseekers were not guided to apply to specific jobs or search in specific areas. This limits the scope for competition between treated and control jobseekers in the experiment.

³²Information about skill comparative advantage as we define it may generate search congestion if communication-heavy jobs are substantially more common than numeracy-heavy jobs or vice versa. However, we showed in Section 3.4 that the two types of jobs are roughly equally common on SAYouth.mobi.

8 Conclusion

We provide evidence of a relatively understudied type of job search friction: misdirected search due to limited information about skill comparative advantage. Giving jobseekers additional information about their skill comparative advantage can shift their beliefs, redirect search effort toward jobs that better match their comparative advantage, and allow them to get better-quality jobs. These problems might be particularly important when jobseekers struggle to evaluate their fit with different jobs when navigating new or changing labor markets due to industrial displacement, migration, or structural transformation ([Huckfeldt, 2022](#); [Robinson, 2018](#)).

As we discuss in the introduction, these findings can help to interpret existing research on jobseekers' beliefs about their labor market prospects, on firm-side limited information about jobseekers' skills, and on skill (mis)match between jobs and workers.

Our approach highlights the benefits of combining multiple experimental designs with multiple measurement strategies. This allows us to observe shifts in beliefs and search direction in a simple two-skill setting with more precise measurement and to study effects on labor market outcomes in a larger sample with a longer timeframe.

We end on some deliberately speculative questions for future research. Our results show private gains for jobseekers who acquire more information about their skill comparative advantage. Can this type of information be efficiently provided to jobseekers outside of research studies? And by whom? We show in [Appendix G](#) that the treatment effect on earnings implies that the average treated jobseeker earns enough extra in the 3.5 months between treatment and endline to cover 1.8 times the average variable cost of the assessment operation. This suggests the possibility of profitable market provision. Assessment allows substantial scale economies, particularly if run on online job search and matching platforms. Many platforms, including SAYouth.mobi, already offer skill assessments to jobseekers ([LinkedIn, 2023](#)). However, prospective private skill assessors might face large fixed costs of developing assessments and building brand credibility, and jobseekers with firmly-held beliefs about their skills might not pay for assessments, even if these beliefs are inaccurate. A strong education system can in principle provide graduates with reliable information about their comparative advantage, reducing the need for market-based provision. But the accuracy of information acquired during schooling may decrease over time, both as people age and as the labor market evolves. Government-funded job search counseling services sometimes include skill assessments but researchers have not yet studied the role of this specific component of job search counseling ([McCall et al., 2016](#)). Future work might examine the economics of both public- and private-sector actors providing information about comparative advantage across different types of skills.

References

- Abebe, G., A. S. Caria & E. Ortiz-Ospina (2021a) "The Selection of Talent: Experimental and Structural Evidence from Ethiopia," *American Economic Review*, 111 (6), 1757–1806.
- Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin & S. Quinn (2021b) "Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City," *Review of Economic Studies*, 88 (3), 1279–1310.
- Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin, S. Quinn & F. Shilpi (2022) "Matching Frictions and Distorted Beliefs: Evidence from a Job Fair Experiment," Working Paper.
- Abel, M., R. Burger & P. Piraino (2020) "The Value of Reference Letters: Experimental Evidence from South Africa," *American Economic Journal: Applied Economics*, 12 (3), 40–71.
- Acemoglu, D. & D. Autor (2011) "Skills, Tasks and Technologies: Implications for Employment and Earnings," in Card, D. & O. Ashenfelter eds. *Handbook of Labor Economics*, 1043–1171: Elsevier.
- Adams-Prassl, A. & A. Andrew (2023) "Revealed Beliefs and the Marriage Market Return to Education," Working Paper.
- Adams-Prassl, A., T. Boneva, M. Golin & C. Rauh (2023) "Perceived Returns to Job Search," *Labour Economics*, 80, 102307.
- Alfonsi, L., O. Bandiera, V. Bassi, R. Burgess, I. Rasul, M. Sulaiman & A. Vitali (2020) "Tackling youth unemployment: Evidence from a labor market experiment in Uganda," *Econometrica*, 88 (6), 2369–2414.
- Alfonsi, L., M. Namubiru & S. Spaziani (2022) "Meet Your Future: Experimental Evidence on the Labor Market Effects of Mentors," Working Paper.
- Algan, Y., B. Crépon & D. Glover (2022) "Are Active Labor Market Policies Directed at Firms Effective? Evidence from a Randomized Evaluation with Local Employment Agencies," Working Paper.
- Altmann, S., A. Falk, S. Jäger & F. Zimmermann (2018) "Learning About Job Search: A Field Experiment With Job Seekers in Germany," *Journal of Public Economics*, 164, 33–49.
- Altmann, S., A. Glenn, R. Mahlstedt & A. Sebald (2022) "The Direct and Indirect Effects of Online Job Search Advice," IZA Discussion Paper 15830.
- Altonji, J., P. Arcidiacono & A. Maurel (2016) "Chapter 7 - The Analysis of Field Choice in College and Graduate School: Determinants and Wage Effects," in Hanushek, E. A., S. Machin & L. Woessmann eds. *Handbook of the Economics of Education*, 5, 305–396: Elsevier.
- Altonji, J. & C. Pierret (2001) "Employer Learning and Statistical Discrimination," *The Quarterly Journal of Economics*, 116 (1), 313–335.
- Anderson, M. L. (2008) "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*, 103 (484), 1481–1495.
- Arcidiacono, P., P. Bayer & A. Hizmo (2010) "Beyond Signaling and Human Capital: Education and the Revelation of Ability," *American Economic Journal: Applied Economics*, 2 (4), 76–104.
- Baley, I., A. Figueiredo & R. Ulbricht (2022) "Mismatch Cycles," *Journal of Political Economy*, 130 (11).
- Bandiera, O., V. Bassi, R. Burgess, I. Rasul, M. Sulaiman & A. Vitali (2023) "The Search for Good Jobs: Evidence from a Six-Year Field Experiment in Uganda," *Journal of Labor Economics*, forthcoming.
- Bandiera, O., A. Elsayed, A. Smurra & C. Zipfel (2022a) "Young Adults and Labor Markets in Africa," *Journal of Economic Perspectives*, 36 (1), 81–100.
- Bandiera, O., N. Parekh, B. Petrongolo & M. Rao (2022b) "Men Are from Mars, and Women Too:

- A Bayesian Meta-analysis of Overconfidence Experiments," *Economica*, 89 (S1), 38–70.
- Banerjee, A. & S. Sequeira (2023) "Learning by Searching: Spatial Mismatches and Imperfect Information in Southern Labor Markets," *Journal of Development Economics*, 164, 103111.
- Banerjee, A. V. & A. F. Newman (1993) "Occupational choice and the process of development," *Journal of Political Economy*, 101 (2), 274–298.
- Bardhan, P. & C. Udry (1999) *Development Microeconomics*: Oxford University Press.
- Bassi, V. & A. Nansamba (2022) "Screening and Signalling Non-Cognitive Skills: Experimental Evidence from Uganda," *The Economic Journal*, 132 (642), 471–511.
- Bazzi, S., L. Cameron, S. G. Schaner & F. Witoelar (2021) "Information, Intermediaries, and International Migration," NBER Working Paper 29588.
- Beam, E. (2016) "Do Job Fairs Matter? Experimental Evidence from the Philippines," *Journal of Development Economics*, 120, 32–40.
- Belot, M., P. Kircher & P. Muller (2019) "Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice," *Review of Economic Studies*, 86 (4), 1411–1447.
- Belot, M., P. Kircher & P. Mueller (2022) "Do the Long-term Unemployed Benefit from Automated Occupational Advice during Online Job Search?," IZA Discussion Paper 15452.
- Benjamini, Y., A. Krieger & D. Yekutieli (2006) "Adaptive Linear Step-Up Procedures That Control the False Discovery Rate," *Biometrika*, 93 (3), 491–507.
- Berry, J., R. Dizon-Ross & M. Jagnani (2022) "Not Playing Favorites: Parents Equalize Their Children's Opportunities," NBER Working Paper 26732.
- Bobba, M. & V. Frisanchio (2022) "Self-Perceptions About Academic achievement: Evidence From Mexico City," *Journal of Econometrics*, 231 (1), 58–73.
- Boudreau, L., R. Heath & T. McCormick (2023) "Migrants, Experience, and Working Conditions in Bangladeshi Garment Factories," *Journal of Economic Behavior and Organization*, forthcoming.
- Card, D., J. Kluve & A. Weber (2018) "What Works? A Meta-Analysis of Recent Active Labor Market Program Evaluations," *Journal of the European Economic Association*, 16 (3), 894–931.
- Carranza, E., R. Garlick, K. Orkin & N. Rankin (2022) "Job Search and Hiring with Limited Information about Workseekers' Skills," *American Economic Review*, 112 (11), 3547–83.
- Chakravorty, B., A. Y. Bhatiya, C. Imbert, M. Lohnert, P. Panda & R. Rathelot (2023) "Impact of the COVID-19 Crisis on India's Rural Youth: Evidence From a Panel Survey and an Experiment," *World Development*, 168, 106242.
- Chen, J. & J. Roth (2023) "Logs With Zeros? Some Problems and Solutions," Working Paper.
- Clark, D. & P. Martorell (2014) "The Signaling Value of a High School Diploma," *Journal of Political Economy*, 122 (2), 282–318.
- Conlon, J. J., L. Pilossoph, M. Wiswall & B. Zafar (2018) "Labor Market Search With Imperfect Information and Learning," NBER Working Paper 24988.
- Cortés, P., J. Pan, L. Pilossoph & B. Zafar (2023) "Gender Differences in Job Search and the Earnings Gap: Evidence from the Field and Lab," *The Quarterly Journal of Economics*, forthcoming.
- Crépon, B., E. Duflo, M. Gurgand, R. Rathelot & P. Zamora (2013) "Do Labor Market Policies Have Displacement Effects? Evidence from a Clustered Randomized Experiment," *The Quarterly Journal of Economics*, 128 (2), 531–580.
- Delavande, A. (2023) "Expectations in Development Economics," in Bachmann, R., G. Topa & W. van der Klaauw eds. *Handbook of Economic Expectations*, Chap. 9, 261–291: Academic Press.
- Diamond, A. (2013) "Executive Functions," *Annual Review of Psychology*, 64, 135–168.
- Dizon-Ross, R. (2019) "Parents' Beliefs about Their Children's Academic Ability: Implications for Educational Investments," *American Economic Review*, 109 (8), 2728–2765.
- Donovan, K., J. Lu & T. Schoellman (2023) "Labor Market Dynamics and Development," *The Quar-*

- terly *Journal of Economics*, forthcoming.
- Duckworth, A. (2016) *Grit: The Power of Passion and Perseverance*, New York, NY, US: Scribner/Simon & Schuster.
- Eil, D. & J. M. Rao (2011) "The Good News-Bad News Effect: Asymmetric Processing of Objective Information About Yourself," *American Economic Journal: Microeconomics*, 3 (2), 114–38.
- Falk, A., D. Huffman & U. Sunde (2006) "Self-Confidence and Search," IZA Discussion Paper 2525.
- Feld, B., A. Osman & A. Nagy (2022) "What Do Jobseekers Want? Comparing Methods to Estimate Reservation Wages and the Value of Job Attributes," *Journal of Development Economics*, 159.
- Fernando, N., N. Singh & G. Tourek (2023) "Hiring Frictions and the Promise of Online Job Portals: Evidence from India," *American Economic Review: Insights*, forthcoming.
- Field, E., R. Garlick, N. Subramanian & K. Vyborny (2023) "Why Don't Jobseekers Search More? Barriers and Returns to Search on a Job Matching Platform," Working Paper.
- Fredriksson, P., L. Hensvik & O. N. Skans (2018) "Mismatch of Talent: Evidence on Match Quality, Entry Wages, and Job Mobility," *American Economic Review*, 108 (11), 3303–3338.
- Garlick, R., K. Orkin & S. Quinn (2020) "Call Me Maybe: Experimental Evidence on Frequency and Medium Effects in Microenterprise Surveys," *The World Bank Economic Review*, 34 (2), 418–443.
- Gibbons, R., L. Katz, T. Lemieux & D. Parent (2005) "Comparative Advantage, Learning, and Sectoral Wage Determination," *Journal of Labor Economics*, 23 (4), 681–724.
- Gneezy, U., A. Rustichini & A. Vostroknutov (2010) "Experience and Insight in The Race Game," *Journal of Economic Behavior and Organization*, 75 (2), 144–155.
- Goldin, C. (1990) *Understanding the Gender Gap*: Oxford University Press.
- Groh, M., D. McKenzie, N. Shammout & T. Viswanath (2015) "Testing the Importance of Search Frictions and Matching through a Randomized Experiment in Jordan," *IZA Journal of Labor Economics*, 4 (7).
- Guvenen, F., B. Kuruscu, S. Tanaka & D. Wiczer (2020) "Multidimensional Skill Mismatch," *American Economic Journal: Macroeconomics*, 12 (1), 210–244.
- Haaland, I., C. Roth & J. Wohlfart (2023) "Designing Information Provision Experiments," *Journal of Economic Literature*, 61 (1), 3–40.
- Hardy, M. & J. McCasland (2023) "Are Small Firms Labor Constrained? Experimental Evidence from Ghana," *American Economic Journal: Applied Economics*, 15 (2), 253–284.
- Harrison, G. & J. List (2004) "Field Experiments," *Journal of Economic Literature*, 42, 1009–1055.
- Heath, R. (2018) "Why Do Firms Hire Using Referrals? Evidence from Bangladeshi Garment Factories," *Journal of Political Economy*, 126 (4), 1691–1746.
- Hensel, L., T. G. Tekleselassie & M. Witte (2022) "Formalized Employee Search and Labor Demand," Working Paper.
- Hoffman, M. & S. Burks (2020) "Worker Overconfidence: Field Evidence and Implications for Employee Turnover and Firm Profits," *Quantitative Economics*, 11, 315–348.
- Huckfeldt, C. (2022) "Understanding the Scarring Effect of Recessions," *American Economic Review*, 112 (4), 1273–1310.
- Huffman, D., C. Raymond & J. Shvets (2022) "Persistent Overconfidence and Biased Memory: Evidence from Managers," *American Economic Review*, 112 (10), 3141–75.
- ILO (2022) "International Labour Organization's Global Employment Trends for Youth 2022," Geneva.
- Ioannides, Y. M. & L. D. Loury (2004) "Job Information Networks, Neighborhood Effects, and Inequality," *Journal of Economic Literature*, 42 (4), 1056–1093.
- Jäger, S., C. Roth, N. Roussille & B. Schoefer (2023) "Worker Beliefs About Outside Options," Working Paper.

- Jepsen, C., P. Mueser & K. Troske (2016) "Labor Market Returns to the GED Using Regression Discontinuity Analysis," *Journal of Political Economy*, 124 (3), 621–649.
- Johnston, A. C. & A. Mas (2018) "Potential Unemployment Insurance Duration and Labor Supply: The Individual and Market-Level Response to a Benefit Cut," *Journal of Political Economy*, 126, 2480–2522.
- Jones, S. & R. Santos (2022) "Can Information Correct Optimistic Wage Expectations? Evidence From Mozambican Job-seekers," *Journal of Development Economics*, 159, 102–987.
- Jovanovic, B. (1979) "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87 (5), 972–990.
- Kahn, L. & F. Lange (2014) "Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures," *Review of Economic Studies*, 84 (1), 1575–1613.
- Kelley, E. M., C. Ksoll & J. Magruder (2023) "How do Digital Platforms Affect Employment and Job Search? Evidence from India ," *Journal of Development Economics*, forthcoming.
- Kerr, A. (2017) "Tax(i)ing The Poor? Commuting Costs in South African Cities," *South African Journal of Economics*, 85 (3), 321–340.
- Kessler, J. B., C. Low & C. D. Sullivan (2019) "Incentivized Resume Rating: Eliciting Employer Preferences without Deception," *American Economic Review*, 109 (11), 3713–44.
- LaLive, R., C. Landais & J. Zweimuller (2022) "Market Externalities of Large Unemployment Insurance Extension Programs," *American Economic Review*, 105, 3564–3596.
- Lam, D., C. Ardington & M. Leibbrandt (2011) "Schooling as a Lottery: Racial Differences in School Advancement in Urban South Africa," *Journal of Development Economics*, 95 (2), 121–136.
- Lentz, R., J. Maibom & E. Moen (2022) "Competitive or Random Search?," Working Paper.
- LinkedIn (2023) "LinkedIn Skill Assessments," <https://www.linkedin.com/help/linkedin/answer/a507734>, Date accessed: 2023/04/21.
- Lise, J. & F. Postel-Vinay (2020) "Multidimensional Skills, Sorting, and Human Capital Accumulation," *American Economic Review*, 110 (8), 2328–76.
- MacLeod, W. B., E. Riehl, J. Saavedra & M. Urquiola (2017) "The Big Sort: College Reputation and Labor Market Outcomes," *American Economic Journal: Applied Economics*, 9 (3), 223–261.
- Malmendier, U. & G. Tate (2015) "Behavioral CEOs: On the Role of Managerial Overconfidence," *Journal of Economic Perspectives*, 29 (4), 37–60.
- Mas, A. & A. Pallais (2017) "Valuing Alternative Work Arrangements," *American Economic Review*, 107 (12), 3722–3759.
- McCall, B., J. Smith & C. Wunsch (2016) "Government-sponsored Vocational Education for Adults," in Hanushek, E., S. Machin & L. Woessmann eds. *Handbook of the Economics of Education*, 5, 479–652: Elsevier.
- McKenzie, D. (2017) "How Effective Are Active Labor Market Policies in Developing Countries? A Critical Review of Recent Evidence," *World Bank Research Observer*, 32 (2), 127–154.
- Moore, D. A. & P. J. Healy (2008) "The Trouble With Overconfidence.," *Psychological Review*, 115 (2), 502.
- Mortenson, D. & C. Pissarides (1994) "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61, 397–415.
- Mueller, A. I. & J. Spinnewijn (2023) "Expectations Data, Labor Market, and Job Search," in Bachmann, R., G. Topa & W. van der Klaauw eds. *Handbook of Economic Expectations*, Chap. 22, 677–713: Academic Press.
- Mueller, A. I., J. Spinnewijn & G. Topa (2021) "Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence and Bias," *Econometrica*, 111, 324–363.
- Mullahy, J. & E. C. Norton (2022) "Why Transform Y? A Critical Assessment of Dependent-

- Variable Transformations in Regression Models for Skewed and Sometimes-Zero Outcomes," NBER Working Paper 30735.
- Nielsen, E. (2023) "How Sensitive are Standard Statistics to the Choice of Scale?", Working Paper.
- Pallais, A. (2014) "Inefficient Hiring in Entry-Level Labor Markets," *American Economic Review*, 104 (11), 3565–3599.
- Pillay, A. L. (2020) "Prioritising Career Guidance and Development Services in Post-Apartheid South Africa," *African Journal of Career Development*, 2 (1), 1–5.
- Posner, M. & G. DiGirolamo (1998) "Executive Attention: Conflict, Target Detection, and Cognitive Control," in Parasuraman, R. ed. *The Attentive Brain*, 401–423: MIT Press.
- Pritchett, L. (2013) *The Rebirth of Education: Schooling Ain't Learning*, Washington, DC: Center for Global Development.
- Raven, J. & J. Raven (2003) "Raven Progressive Matrices," in McCallum, R. ed. *Handbook of Non-verbal Assessment*, 223–237, Boston: Springer.
- Robinson, C. (2018) "Occupational Mobility, Occupation Distance, and Specific Human Capital," *Journal of Human Resources*, 53 (2), 513–551.
- Rosenberg, M. (1965) *Society and the Adolescent Self-Image*: Princeton University Press.
- Roy, A. D. (1951) "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3 (2), 135–146.
- Sanders, C. & C. Taber (2012) "Life-Cycle Wage Growth and Heterogeneous Human Capital," *Annual Review of Economics*, 4 (1), 399–425.
- Santos-Pinto, L. & L. E. de la Rosa (2020) "Overconfidence in Labor Markets," *Handbook of Labor, Human Resources and Population Economics*, 1–42.
- Schoer, V., M. Ntuli, N. Rankin, C. Sebastiao & K. Hunt (2010) "A Blurred Signal? The Usefulness of National Senior Certificate (NSC) Mathematics Marks as Predictors of Academic Performance at University Level," *Perspectives in Education*, 28 (2), 9–18.
- Sockin, J. & A. Sojourner (2023) "What's the Inside Scoop? Challenges in the Supply and Demand for Information about Job Attributes," *Journal of Labor Economics*, forthcoming.
- Spinnewijn, J. (2015) "Unemployed but Optimistic: Optimal Insurance Design With Biased Beliefs," *Journal of the European Economic Association*, 13 (1), 130–167.
- SSA (2016) *Quarterly Labor Force Survey Quarter 4 2016*, Pretoria: Statistics South Africa.
- (2022) *Quarterly Labor Force Survey Quarter 3 2022*, Pretoria: Statistics South Africa.
- Stroop, J. R. (1935) "Studies of Interference in Serial Verbal Reactions," *Journal of Experimental Psychology*, 18 (3), 643–662.
- Subramanian, N. (2022) "Workplace Attributes and Women's Labor Supply Decisions: Evidence from a Randomized Experiment," Working Paper.
- Taylor, T. (1994) "A Review of Three Approaches to Cognitive Assessment, and a Proposed Integrated Approach Based on a Unifying Theoretical Framework," *South African Journal of Psychology*, 24.
- Wheeler, L., R. Garlick, E. Johnson, P. Shaw & M. Gargano (2022) "LinkedIn(to) Job Opportunities: Experimental Evidence from Job Readiness Training," *American Economic Journal: Applied Economics*, 14 (2), 101–25.
- Wiswall, M. & B. Zafar (2015) "Determinants of College Major Choice: Identification Using an Information Experiment," *Review of Economic Studies*, 82 (2), 791–824.
- Wu, L. (2021) "Partially Directed Search in the Labor Market," Working Paper.

Jobseekers’ Beliefs about Comparative Advantage and (Mis)Directed Search: Online Appendices Not for Publication

This online material contains twelve appendices. Appendix A contains appendix summary statistics tables. Appendix B describes our skill measurements and show that firms not only value the skills we study but are also able to detect them. In Appendix C we describe how the skill beliefs were elicited and provide further descriptive statistics on these beliefs. Appendix D describes the experimental protocols in detail. We collect robustness checks and various heterogeneous treatment effect results in Appendix E. Appendix F provides an additional analysis of the role of labor market beliefs in shaping search direction. Appendix G details the cost-benefit calculation. Appendix H contains additional exhibits and results that show that changes in search effort are unlikely to explain our results – as discussed in Section 5. Appendix I contains exhibits supporting the additional mechanism analyses from Section 6. Appendix J shows the results related to gender heterogeneity. Finally, Appendix K describes how our analysis relates to the pre-analysis plans.

A Sample Description

This appendix contains two additional sample description tables. Table A.1 shows online search activity on the SAYouth.mobi by the tight experiment sample in the 30 days prior to treatment. Table A.2 provides summary statistics for jobseekers in the big experiment. Comparing this table to Table 1 confirms that the two sample are very similar.

Table A.1: Summary Statistics for Search on SAYouth.mobi Platform - Tight Experiment

	Mean (1)	Median (2)	Min (3)	Max (4)	SD (5)	Obs. (6)
# days active on platform	3.17	2.00	0.00	25.00	4.06	278
# applications clicks (winsorized)	6.53	2.00	0.00	66.00	11.27	278
# applications clicks for numeracy heavy jobs	0.39	0.00	0.00	11.00	1.10	278
# applications clicks for communication heavy jobs	0.82	0.00	0.00	9.00	1.64	278
Fraction of skill coded application clicks	0.12	0.00	0.00	1.00	0.20	278

Notes: Table A.1 shows summary statistics of participants’ online job search on the SAYouth.mobi platform in the 30 days prior to the intervention in the tight experiment. Application clicks are defined as initiating an application by clicking on the “apply here” button. The platform does not consistently record whether candidates complete applications after starting them. Winsorized variables are winsorized at the 99th percentile.

Table A.2: Summary Statistics - Big Experiment

	(1) Mean	(2) Median	(3) Min	(4) Max	(5) SD	(6) Obs.
<u>Panel A: Demographics</u>						
Black African	0.98	1.00	0.00	1.00	0.12	4389
Male	0.38	0.00	0.00	1.00	0.48	4389
Age	23.67	23.14	18.04	35.08	3.28	4389
Completed secondary education only	0.61	1.00	0.00	1.00	0.49	4389
University degree / diploma	0.17	0.00	0.00	1.00	0.37	4389
Any other post-secondary qualification	0.22	0.00	0.00	1.00	0.41	4389
<u>Panel B: Labor market background</u>						
Any work in last 7 days	0.37	0.00	0.00	1.00	0.48	4389
Has worked in permanent wage job before	0.09	0.00	0.00	1.00	0.29	4377
Earnings in USD (last 7 days, winsorized)	31.26	0.00	0.00	476.00	75.72	4389
<u>Panel C: Search behavior</u>						
Any job search in last 7 days	0.97	1.00	0.00	1.00	0.17	4389
# applications (last 30 days, winsorized)	9.34	5.00	0.00	90.00	12.85	4346
Search expenditure in USD (last 7 days, winsorized)	30.97	20.40	0.00	204.00	33.05	3995
Hours spent searching (last 7 days, winsorized)	17.06	8.00	0.00	96.00	19.68	4273
# job offers (last 30 days, winsorized)	0.80	0.00	0.00	20.00	2.66	4335
<u>Panel D: Skills beliefs</u>						
Aligned belief about CA	0.20	0.00	0.00	1.00	0.40	4312
Fraction of aligned belief domains	0.38	0.33	0.00	1.00	0.31	4378

Notes: **Table A.2** shows summary statistics for the big experiment. “CA” stands for comparative advantage. Winsorized variables are winsorized at the 99th percentile. All monetary values are in 2021 USD purchasing power parity terms.

B Skill Assessments and Firms’ Valuation of Skills


This appendix describes the skill assessments, assessment results, and evidence that firms both value the skills we study and have some information about job applicants’ skill levels. Information on participant recruitment, workshop structure, and the timing of the assessments relative to treatment and surveys is described in Appendix D.

B.1 Measurement

Tight Experiment: We assess jobseekers’ skills on numeracy, communication, and concept formation. The *numeracy* assessment was based on a test developed by a retail chain in South Africa. The chain uses the test to assess candidates’ skills needed to become a cashier. The *communication* assessment quantifies jobseekers’ listening and reading comprehension. The assessment focuses on high school English proficiency and was developed by a local adult education provider (www.mediaworks.co.za). The *concept formation* assessment was a test similar to Raven’s matrices (Raven & Raven, 2003). This assesses jobseekers’ fluid intelligence: their conceptual reasoning and the rate at which they learn.

We find ample variation within skills and across skill quintiles. Figure B.1 shows the distribution of standardized numeracy and communication scores in the tight experiment. The distributions show no evidence of floor or ceiling effects. Table B.1 shows the joint distribution of numeracy skill quintiles and communication skill quintiles. Table B.2 shows the correlations between skill measures in the tight experiment. All three skill measures have pairwise correlations > 0.3 in the full sample. Communication and numeracy have zero correlation in the restricted sample, which excludes jobseekers with tied communication and numeracy quintiles, because this restriction mechanically removes some of the positive relationship between the skills.

Figure B.1: Distribution of Standardized Assessment Scores - Tight Experiment



Notes: Figure B.1 shows the distribution of standardized assessment scores in the tight experiment. Panel A displays the fraction of standardized numeracy scores in bins. Panel B displays the same for communication scores.

Table B.1: Joint Distribution of Skill Quintiles - Tight Experiment

Numeracy quintile	Communication quintile				
	Bottom	Lower middle	Middle	Upper middle	Top
Bottom	0.00%	7.19%	6.47%	4.68%	4.68%
Lower middle	11.51%	0.00%	7.19%	7.91%	9.35%
Middle	3.96%	2.52%	0.00%	4.68%	2.88%
Upper middle	2.52%	4.32%	3.96%	0.00%	7.19%
Top	0.72%	2.16%	0.72%	5.40%	0.00%

Notes: Table B.1 shows that there considerable variation in quintiles across numeracy and communication skills in the tight experiment. The sample is restricted to the 278 jobseekers with a unique comparative advantage in skills. This forces entries on the main diagonal to be zero.

Table B.2: Correlation between Skill Quintiles - Tight Experiment

	Numeracy (1)	Communication (2)	Concept formation (3)
<u>Panel A: Restricted sample</u>			
Numeracy	1.000	-0.004	0.284
Communication		1.000	0.216
Concept formation			1.000
<u>Panel B: Full sample</u>			
Numeracy	1.000	0.306	0.375
Communication		1.000	0.326
Concept formation			1.000

Notes: Table B.2 shows that correlation coefficients between the skills in the tight experiment are mostly moderate-sized and positive. Panel A shows results for the sample with a unique comparative advantage between numeracy and communication. Therefore, the correlation between those two skills is zero in this subsample. Panel B shows results for the full sample.

Big Experiment: We measured six skills: communication, numeracy, concept formation, focus, grit, and planning. The *communication* assessment is equivalent to the assessment used in the tight experiment. It measures listening and reading comprehension skills. The *numeracy* assessment measures practical arithmetic skills and pattern recognition. The first part of the assessment is the same as the numeracy assessment in tight experiment. This part assesses jobseekers' skills needed for a cashier position. The second part of the numeracy assessment was developed by an adult education provider (www.mediaworks.co.za) and measures high school mathematics skills including calculations involving money, time, areas and quantities. For measuring *concept formation* skills, we used a test that is very similar to Ravens' matrices (Raven & Raven, 2003). The *focus* measure is a computerized, color-based Stroop task (Stroop, 1935). It evaluates jobseekers' inhibitory control, controlling one's attention and guiding thought and action to achieve a goal (Diamond, 2013; Posner & DiGirolamo, 1998). We assess jobseekers' *grit* using the self-reported 8-item scale from Duckworth (2016). It rates jobseekers' willingness to work on difficult tasks and persevere to achieve long-term goals. Finally, *planning* measures how jobseekers are able to search for relevant information and anticipate the consequences of actions. The assessment adapts the Hit 15 task in Gneezy et al. (2010), in which the computer and the participant take turns to add one, two, or three points to a point basket. The party whose action leads to a score of 15 in the point basket wins.

Table B.3 shows correlations between terciles of skills in the big experiment. Scores are weakly positively correlated across assessments, with pairwise correlations of 0.05–0.44. Hence, the assessments horizontally differentiate candidates based on their relative skills rather than only ranking or vertically differentiating them on a single skill dimension.

Table B.3: Correlation Between Skill Terciles - Big Experiment

	Concept formation (1)	Communication (2)	Numeracy (3)	Grit (4)	Planning (5)	Focus (6)
Concept formation	1.000	0.298	0.435	0.098	0.214	0.189
Communication		1.000	0.331	0.095	0.213	0.173
Numeracy			1.000	0.139	0.266	0.159
Grit				1.000	0.090	0.047
Planning					1.000	0.173
Focus						1.000

Notes: Table B.3 shows that correlation coefficients between the skills in the big experiment are mostly moderate-sized and positive.

B.2 Firms Value These Skills

In this section we provide support that the skills we measured are important in this labor market.

First, the communication, concept formation, and numeracy assessments have been used to screen jobseekers by our partner, the [Harambee Youth Employment Accelerator](#). By 2016, Harambee had been contracted by firms in South Africa to screen roughly 160,000 prospective workers using these assessments. This suggests a revealed willingness of firms to pay to learn the results of these assessments. However, this does not mean that assessment results are the only information firms use in their hiring decisions. We do not assume that firms use the information at their disposal optimally, and thus, we do not claim that these tests are the best predictors of jobseekers' productivity.

Second, we used an incentivized choice experiment to show that firms vary in their valuation of communication and numeracy skills and value both highly relative to some forms of education. For this data collection, we recruited 67 firms soon after the big experiment by going door-to-door in areas of Johannesburg where most of the jobseekers in the big experiment lived. 81% of firms are in the retail or hospitality sectors, where many jobseekers in both experiments applied for jobs. Recruited firms have a mean size of 15 workers, half of whom are in entry-level roles, and planned to hire an average of 4 new entry-level workers in the next year.

Importantly for our current argument, we measured the preferences of these firms over the six skills used in the big experiment relative to each other and to additional education. Each firm was asked to rank multiple jobseeker profiles with different levels of skills and with or without a one-year post-secondary diploma, all with completed secondary school. To incentivize the choices, we used firms' rankings to match them with jobseekers with specific skill profiles from Harambee's database, in a similar spirit to

Kessler et al. (2019).

Table B.4 shows the average ranking of numeracy, communication and education over the 67 firms. There are six different possible rankings of these three elements, each shown in a row. The shares of firms in these bins are shown in column 4. Column 6 collapses these shares based on the most important skill. In this column, we see that 57% of firms prefer a candidate with top-tercile numeracy skills, 34% prefer a candidate with top-tercile communication skills and only 9% firms prefer a candidate with a relatively better educational achievement, i.e. a candidate with a diploma but with only middle-tercile communication and numeracy skills.

Table B.4: Firms' Preference Ranking Over Communication Skills, Numeracy Skills, and Formal Education

	Top (1)	Middle (2)	Bottom (3)	Share (%) (4)	Most Important Skill (5)	Share (%) (6)
1	Num	Comm	Educ	52.24	Numeracy	56.72
2	Num	Educ	Comm	4.48		
3	Comm	Num	Educ	28.36	Communication	34.33
4	Comm	Educ	Num	5.97		
5	Educ	Num	Comm	1.49	Education	8.96
6	Educ	Comm	Num	7.46		

Notes: Table B.4 shows that firms, on average, value applicants' numeracy (Num) and communication skills (Comm) more than a one-year post-secondary certificate (Educ) and vary in their relative ranking of communication and numeracy skills. The results are based on an incentivized choice experiment with 67 small and medium sized businesses in Johannesburg. The rows represent all the possible rankings. Col. 4 shows the share of firms who chose the respective rankings. Cols. 5 and 6 do the same but collapse the rows according to the most important skill.

B.3 Observability of Skills and Firms' Valuation of Applicant Skill Match

As part of the tight experiment, we conducted a measurement exercise to show that firms partially observe assessed skills and value applicants whose skill profile matches their job requirements. During the job search workshop, we asked jobseekers to choose between applying for a real communication- or numeracy-heavy job at a firm that hires for a range of entry-level roles, including call center and data capture jobs. (See Appendix D.1 for details about the task.) Jobseekers prepared a CV and a cover letter during the workshop, both designed for general use rather than tailored to these specific jobs. Two members of the firm's HR team evaluated every applicant for both jobs based on their CV and generic cover letter. Evaluators did not know which applicant applied for which job and

Table B.5: Employer Evaluation of Job Applicants Based on Skills

	Mean			Difference			Obs.
	Aligned	Non-aligned	SD	Δ	Δ/SD	$p(\Delta = 0)$	
	(1)	(2)	(3)	(4)	(5)	(6)	
<u>Panel A: Skill levels</u>							
Skill (1-5)	2.93	2.79	0.66	0.15	0.22	0.01	277
<u>Panel B: Job-related evaluation</u>							
Overall score (1-5)	3.00	2.84	0.86	0.16	0.18	0.05	277
Interview invitation (dummy)	0.43	0.34	0.49	0.09	0.18	0.07	277

Notes: Table B.5 shows that the HR team of an employer can observe jobseekers' skills and evaluates applicants more highly if their assessed comparative advantage in skills matches the job's requirements. One pair of the job choice task advertisements was from a firm that hires for a range of entry-level roles. We submitted the jobseekers' coverletter and resume to the firm based on which two members of the firm's HR team evaluated every applicant for both jobs. Evaluators rated the jobseekers' skill levels (Panel A) as well as their general suitability for the job (on a scale from 1 to 5) and whether they would invite the candidate for an interview (Panel B). We show the mean outcomes across evaluators for the skill/ job that is aligned with the jobseekers' comparative advantage in col. 1; and the outcomes for the misaligned skill / job in col. 2. The pooled standard deviation of the measures in col. 1 and 2 are in col. 3. Col. 4 shows the difference between cols. 1 and 2. Col. 5 shows this difference in terms of standard deviations. Col. 6 shows the p-value associated with a test of equality across cols. 1 and 2. Col. 7 shows the number of observations.

were not shown applicants' skill assessment results. We received data on the evaluators' assessment of each applicant's communication skills, numeracy skills, suitability for each type of job, and whether they were recommended for an interview for each type of job.

This measurement exercise shows that the firm's HR team can get some information about jobseekers' skill levels from their application materials. Table B.5, panel A, column 5 shows that the HR team's assessments of skills are positively but not perfectly correlated with our measures of skill: they assigned a 0.22 standard deviation higher score to the skill dimension in which the jobseeker had a comparative advantage on our assessments ($p = 0.01$). The HR team's applicant ratings are also correlated with our measures of skill: they rate the applicant as 0.18 standard deviations more suitable for the job aligned with that applicant's comparative advantage (panel B, row 1, column 6). HR managers were also 9 percentage points more likely to recommend interviewing the candidate for the job aligned with that applicant's comparative advantage (panel B, row 2, column 4). This is a 21% increase relative to a 43% interview recommendation rate in the non-aligned job. These patterns show that skills are at least partly observable to the firm even when jobseekers could not tailor their resumes or cover letters to the specific role. Observability might be greater in natural job search where jobseekers can tailor their applications.

The evidence that jobseekers have different skills, that firms value these skills but differ in which skill they value more, and that firms can at least partly observe skills suggests that redirecting jobseekers' search towards jobs that match their comparative advantage in skills has the potential to improve their labor market outcomes.

C Skill Beliefs

C.1 Skill Beliefs Measurement

This appendix describes and justifies our choices of our main belief measures and explains how we measure beliefs. Our empirical results are very robust to the choice of belief measure. However, there are slight conceptual differences in measurement between and within experiments that we clarify in this appendix.

We measure skill beliefs at the skill $\times$ individual-level in terms of quintiles (tight experiment) or terciles (big experiment). In the tight experiment, we measure beliefs about numeracy and communication skills. In the big experiment, we measure beliefs about numeracy, communication, and concept formation skills. See Appendix B for a description of the assessments. Further, we measure two types of skill beliefs: beliefs about *assessment results* and beliefs about *general* skills. Table C.1 contains the exact wording of our skill belief elicitation questions for both. In practice, these two belief measures are strongly positively correlated within skill, suggesting that jobseekers view the assessments as relevant to their general skills.

Before eliciting beliefs, we define the skills in the following way: "Numeracy means working with numbers. It includes using addition, subtraction, multiplication, and division to solve real problems involving money, time, and quantities. For example, if a box holds 18 cans of tuna, can you calculate how many cans of tuna there are in 9 boxes? Communication means reading, writing, and listening in English. It includes understanding your coworkers and customers when they explain problems they have and explaining how to solve these problems. These are not skills about how to treat other people, just English skills."

Beliefs about General Skills: We measure each jobseeker's beliefs about their general skills as their beliefs about their skills relative to the skills of the reference group in a specific domain, abstracting from specific assessment results. We see these skill beliefs as being *most relevant for search decisions* because they capture general, labor market relevant skills that are not affected by the idiosyncrasy of the performance of our assessments. Put differently, what matters to employers is not how well one does on a specific assessment but rather how well one is able to use a skill consistently at work relative to others. Thus, we use these beliefs to define our preferred measure of aligned comparative advantage

beliefs.

In the tight experiment, we measure beliefs about general skills before and after treatment. (See Figure D.1 that summarizes the experimental design.) We use general skill beliefs in the tight experiment for the descriptive statistics in the summary Table 1. We use general skill beliefs to define our main belief outcome measures in the tight experiment (Table 2) and the corresponding appendix tables (Tables C.2, C.8, E.1, E.7, H.1 (columns 1 to 4), and J.3).

In the big experiment, we only measure beliefs about general skills for a random subsample of participants after treatment. As a robustness check, we estimate treatment effects using the comparative advantage beliefs defined using general skills. We find that treatment increases aligned comparative advantage beliefs by 7.6 pp ($p = 0.003$) and 10.8pp ($p < 0.001$). These are smaller than the effects on beliefs about assessment results but qualitatively similar and still highly significant.

Beliefs about Assessment Results: We measure jobseekers' beliefs about their relative placement in the assessments they completed in both experiments. We consider these beliefs as a *proxy for the information content of the interventions* for each individual because the treatment provides information about relative assessment results (following Haaland et al., 2023). Jobseekers who have inaccurate beliefs about their relative assessment results after taking the assessment will learn that their actual performance differed from their beliefs. Conversely, jobseekers with initially accurate beliefs should not update their beliefs about their assessment results (though they might still become more certain about their beliefs). We hypothesize that individuals with initially accurate beliefs about their comparative advantage in the assessments should have smaller treatment effects. Hence, we estimate heterogeneity using a dummy variable indicating accurate baseline beliefs about jobseekers' comparative advantage in the assessments throughout both the tight and big experiment (equation 4).

In the tight experiment, we asked about beliefs about assessment results after the assessment but before the treatment for the whole sample (see Figure D.1). We use these measures for all heterogeneity analysis by baseline beliefs in the tight experiment: Tables 2 and 3 and all tables in the appendices that report heterogeneity by skill beliefs in the tight experiment. We also ask the same question again right after the treatment administration for the treatment group only to check whether they understood the results on the report. We report this understanding check on page 18. We did not ask the control group again to avoid asking the same question twice in a short amount of time without providing additional information.

On average, 54% of jobseekers in the tight experiment have aligned comparative ad-

Table C.1: Measurement of beliefs about comparative advantage

Description	Survey question
Panel A: Tight experiment	
Beliefs about general skills, pre-treatment (most likely quintile)	Think about 100 people who are jobseekers from Johannesburg aged 18-34 with a matric from a township or rural school. Imagine that we rank everyone according to their [numeracy/communication] skills, from lowest to highest. We create five equal size groups. The first group are the 20 people with the strongest [numeracy/communication] skills. The second group are the 20 people with the next best skills – they are less good than the top 20, but better than the other 60 people. The fifth group are the 20 people with less strong numeracy skills than the other 80. Out of these five groups we just talked about, what group do you think you are most likely to be in based on your [numeracy/communication] skills?
Beliefs about general skills belief, post treatment (most likely quintile)	Think about 100 people who are jobseekers from Johannesburg aged 18-34 with a matric from a township or rural school. Imagine that we rank everyone according to their [numeracy/communication] skills, from lowest to highest. This ranking is based on overall [numeracy/communication] skills, not only the numeracy skills that were tested in the Numeracy assessment you just took. We create five equal size groups. The first group are the 20 people with the strongest [numeracy/communication] skills. The second group are the 20 people with the next best skills – they are less good than the top 20, but better than the other 60 people. The fifth group are the 20 people with less strong [numeracy/communication] skills than the other 80. Out of these five groups we just talked about, what group do you think you are most likely to be in based on your [numeracy/communication] skills?
Beliefs about assessment results, pre- and post-treatment (most likely quintile)	Think about 100 people who are jobseekers from Johannesburg aged 18-34 with a matric from a township or rural school. Imagine that we rank everyone according to their results on the [numeracy/communication] assessment. We create five equal size groups. The first group are the 20 people with the highest numeracy results. The second group are the 20 people with the next best results – they are less good than the top 20, but better than the other 60 people. The fifth group are the 20 people with lower strong numeracy results than the other 80. Out of these five groups we just talked about, what group do you think you are most likely to be in based on your [numeracy/communication] assessment result?
Panel B: Big experiment	
Beliefs about general skills, post-treatment (most likely tercile)	Remember that people who come to Harambee are from Johannesburg, are aged 18-34 and have a matric from a township or rural school. So that should be the group you're picturing. If we ranked candidates by their [numeracy/communication/concept formation] skills, do you think you are in the top third, middle third or bottom third of Harambee candidates?
Beliefs about assessment results, pre-treatment (most likely tercile)	Now think about all the people who are in the room with you. They are all jobseekers from Johannesburg aged 18-34 with a matric from a township or rural school and have done the Harambee assessments. Imagine we line everyone up according to what score they got, from lowest to highest. Then we divide the group into three. The lower third are the people who got the lowest scores. The top third are the people who got the highest scores. The middle third are the rest of the people. Would you be in the top third, middle third or bottom third of people on the [numeracy/communication/concept formation] test?
Beliefs about assessment results, post-treatment (most likely tercile)	Do you remember the assessments you took at Harambee during Phases 1 and 2? [wait for yes]. Now I want you to imagine other Harambee candidates who have also taken these assessments. Remember that people who come to Harambee are from Johannesburg, are aged 18-34 and have a matric from a township or rural school. So that should be the group you're picturing. Imagine we look at everyone's assessment scores, and we make three groups: One group for people with the lowest scores, one group for people with the highest scores, and one group for people in the middle. Each group contains one third of the people who took the assessment. Keep this scenario in your mind, and answer the following questions. Remember that this will not have any impact on your progress with Harambee. These answers are only for research purposes and will be kept confidential. Off the top of your head, do you think you are in the top third, middle third or bottom third of people on the [numeracy/communication/concept formation] test?

Notes: Table C.1 displays the exact wording of our questions to jobseekers about their beliefs about their skills. Note that we construct beliefs about skill comparative advantage from the beliefs about terciles/quintiles of different skills.

vantage beliefs using this measure. The results are similar when we use the measure of general skills for heterogeneity instead of assessment beliefs because the two measures of comparative advantage beliefs are highly correlated ($\rho = 0.68$). Similarly, regressing domain specific general skill beliefs on beliefs about assessment results produces coefficients of 0.39-0.52 across the two skills, with or without controls for assessment results and demographic characteristics.

In the big experiment, we ask both the control and treatment group twice about their beliefs about their assessment results. (See Figure D.2.) First, we ask them at baseline after they took the assessments but before treatment. Second, we ask them again at end-line about 3.5 months after treatment. Given that we only measured general skills for a subsample of jobseekers after treatment, we use beliefs about assessment results both as outcomes and as heterogeneity variables for the big experiment throughout the paper and appendix. However, treatment effects on beliefs about general skills are similar, as we note above.

C.2 Skill Belief Descriptive Statistics and Treatment Effects

Here we provide descriptive statistics about jobseekers' beliefs about skills and their relationship with high school exam results, expectations for job search outcomes, and job search activities. We then present additional treatment effects on beliefs about skills.

Descriptive statistics: Table C.2 displays correlations between prior beliefs about skills and assessment results in the tight experiment. We find that beliefs are strongly correlated with assessment results for numeracy but not for communication. Table C.3 displays correlations between jobseekers' self-reported scores in English and mathematics in the high school leaving exam (matric) and their assessment results and beliefs. Columns 1-3 show that the matric results correlate with our assessments and the score differences between the respective subjects on those exams positively correlate with the measured comparative advantage in the tight experiment, although the latter relationship is imprecisely estimated.

We also find that the high school leaving exam results predict jobseekers' beliefs (Table C.3, columns 4-7). However, we do not find that any baseline demographic variables, including labor market exposure, meaningfully predict having aligned comparative advantage beliefs at baseline in the tight experiment (Table C.4). Table C.5 shows that beliefs are not updated over time in the control group in the big experiment. Overall, this set of results suggests that jobseekers have limited scope to learn about their comparative advantage in this labor market.

Finally, we show additional evidence that jobseekers' beliefs about skills are related to

perceptions of the labor market and search direction in the big experiment, using data for only the control group. Table C.6 shows that jobseekers' beliefs about their comparative advantage correlate positively and significantly with search direction in the control group of the big experiment. Table C.7 shows that jobseekers' with higher beliefs about their skills expect shorter search durations and higher wages.

Additional treatment effects: Underconfident skills are more likely to update than overconfident skills. Table C.8 displays treatment effects on dummies indicating the fraction of under and overconfident skill beliefs. We find that the relative reduction in underconfident beliefs (30.6% and 28.2% of the control mean in the tight and big experiment respectively) is significantly larger than the relative reduction in overconfident beliefs (3.5% and 21.9% of the control mean in the tight and big experiment respectively). This aligns with lab evidence on asymmetric updating (Eil & Rao, 2011).

Jobseekers might update their beliefs about their skill comparative advantage through two channels: because they update beliefs about (1) their own skill levels or (2) the distribution of other jobseekers' skills. To evaluate these channels, we ask jobseekers how many questions they answered correctly on each assessment, after they took the assessments but before any treatment. We construct a dummy for **accurate score ranking**, equal to one if and only if the jobseeker correctly identifies the assessment on which she answered more questions correctly. We then identify jobseekers who have an accurate score ranking and a misaligned comparative advantage belief. These jobseekers can only update their comparative advantage beliefs due to channel (2); channel (1) is largely shut down for them. This admittedly small subset of jobseekers substantially update their comparative advantage beliefs ($p = 0.001$). This suggests that learning about the distribution of other jobseekers' skills at least partially explains the updating of comparative advantage beliefs.

Table C.2: Association between Assessed and Believed Skill Quintiles - Tight Experiment

	Skill quintile beliefs	
	Numeracy (1)	Communication (2)
Numeracy quintile	0.188*** (0.049)	-0.011 (0.034)
Communication quintile	-0.061 (0.040)	0.018 (0.032)
Dep var. mean	2.367	3.259
Observations	278	278

Notes: **Table C.2 shows that jobseekers' beliefs about their score quintiles correlate with their assessment results for numeracy but not communication.** No control variables are included. Heteroskedasticity-robust standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.3: Association between High School Graduation Exam Results and Assessed and Believed Skills - Tight Experiment

	Assessed skills			Beliefs about skills			
	Skill quintile		Comp. adv.	Skill quintile		Comp. adv.	
	Num. (1)	Com. (2)	Num. (3)	Num. (4)	Com. (5)	Num. (6)	Com. (7)
Matric: Math score	1.451*** (0.506)	0.594 (0.609)		1.782*** (0.399)	-0.360 (0.261)		
Matric: English score	0.219 (0.400)	1.219** (0.477)		-0.355 (0.335)	1.016*** (0.223)		
Matric: Δ Math score - English score			0.192 (0.146)			0.308*** (0.108)	-0.487*** (0.139)
Dep var. mean	1.540	2.173	0.378	2.367	3.259	0.137	0.579
Observations	263	263	263	263	263	263	263

Notes: **Table C.3 shows that self-reported scores in English and mathematics in the high school graduation exam (matric) correlate positively with jobseekers' assessment results and their baseline beliefs about skills.** Cols. 1, 2, 4, and 5 show the relationship between math and English matric scores and assessed skill quintiles (cols. 1 and 2) and beliefs about skill quintiles (cols. 4 and 5). Matric scores are rescaled to range from 0 to 1. So, for example, the coefficient in column 4 row 1 shows that moving from the lowest to highest possible matric grade in math is associated with a belief that numeracy is 1.78 quintiles higher. No further control variables are included. The sample size is 263 because 15 jobseekers cannot recall their exam results. Heteroskedasticity-robust standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.4: Association Between Baseline Aligned Comparative Advantage Belief and Other Baseline Characteristics - Tight Experiment

	Aligned CA belief	
	(1)	(2)
Age	-0.012 (0.009)	-0.004 (0.008)
Female	0.048 (0.064)	0.024 (0.059)
Has completed secondary education only	-0.008 (0.114)	-0.024 (0.109)
Has a post secondary certificate	-0.059 (0.133)	-0.071 (0.126)
Has a post secondary diploma	0.095 (0.140)	0.058 (0.135)
Has a post secondary degree	-0.208 (0.140)	-0.208 (0.140)
Employed in any form at baseline	0.030 (0.068)	0.026 (0.063)
Total work experience at baseline (years)	0.017 (0.013)	0.013 (0.012)
Numeracy assessment score (%)		-0.007** (0.003)
Communication assessment score (%)		0.012*** (0.002)
Constant	0.739** (0.244)	0.076 (0.283)
# jobseekers	278	278
R2 (not adjusted)	0.031	0.177
p: all coefficients = 0	0.233	0.000

Notes: Table C.4 shows that baseline comparative advantage beliefs are uncorrelated with demographic characteristics. It displays coefficients from regressions with baseline data of an indicator for aligned comparative advantage belief on age, gender, education categories (omitting less than completed high school), employment, and total work experience and, in the second column only, communication and numeracy assessment scores. Robust standard errors shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.5: Development of Skill Beliefs Over Time - Big Experiment Control Group

	Aligned CA belief			% aligned skill belief		
	(1)	(2)	(3)	(4)	(5)	(6)
Endline	-0.004 (0.012)	-0.001 (0.017)	-0.015 (0.014)	0.009 (0.009)	0.009 (0.013)	0.004 (0.011)
Endline $\times$ Above median search effort		-0.008 (0.024)			0.000 (0.019)	
Above median search effort		0.035** (0.017)			0.003 (0.013)	
Endline $\times$ Worked last 7 days			0.034 (0.027)			0.015 (0.020)
Worked last 7 days			-0.005 (0.019)			0.013 (0.014)
Constant	0.200*** (0.008)	0.183*** (0.012)	0.202*** (0.010)	0.379*** (0.006)	0.377*** (0.009)	0.375*** (0.008)
Observations	4405	4315	4315	4456	4365	4365

Notes: Table C.5 shows that the misalignment of skill beliefs persists over time in the control group of the big experiment even for those who are employed or have above-median search activity. Estimation is at the survey round times jobseeker level and is restricted to the control group. “CA” stands for comparative advantage. Cols. 1 to 3 show results for aligned beliefs about comparative advantage. Cols. 4 to 6 show results for the fraction of aligned skill beliefs. Heteroskedasticity-robust standard errors are in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.6: Association between Skill Beliefs and Search Direction - Big Experiment Control Group

	Target numeracy jobs		Target communication jobs	
	(1)	(2)	(3)	(4)
Believed numeracy CA	0.100*** (0.018)	0.091*** (0.018)	-0.094*** (0.021)	-0.084*** (0.022)
Believed communication CA	-0.079*** (0.021)	-0.067*** (0.020)	0.087*** (0.024)	0.088*** (0.023)
Numeracy CA		0.021 (0.018)		-0.012 (0.022)
Communication CA		-0.060*** (0.019)		0.045** (0.023)
Control mean	0.222	0.222	0.471	0.471
Observations	2183	2179	2183	2179
Controls	No	Yes	No	Yes

Notes: Table C.6 shows that jobseekers’ beliefs about comparative advantage correlate positively with skill-directed search in the control group of the big experiment. “CA” stands for comparative advantage. Dependent variables are dummies indicating that jobseekers rate numeracy (columns 1-2) or communication (columns 3-4) as the most important skill for the jobs that they applied to in the last 30 days, measured at endline. Independent variables are dummies assessed and believed skill comparative advantage. Columns 2 and 4 include controls for age, gender, having worked in a wage job, as well as dummies for three education categories, and measured comparative advantage. Heteroskedasticity-robust standard errors are in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.7: Association between Skill Beliefs and Beliefs About Search Outcomes - Big Experiment Control Group

	E[search duration] (months, w)		E[wage] (w)	
	(1)	(2)	(3)	(4)
Average skill tercile belief (z-scored)	-0.127*** (0.047)	-0.129*** (0.046)	38.318*** (9.518)	33.755*** (9.429)
Control mean	2.718	2.718	892.045	892.045
Observations	2148	2144	2183	2179
Controls	No	Yes	No	Yes

Notes: Table C.7 shows that jobseekers' beliefs about their skills correlate positively with their beliefs about the returns to search in the control group in the big experiment. All specifications control for average standardized skill levels. Columns 2 and 4 further include controls for age, gender, having worked in a wage job, as well as dummies for three education categories. Winsorized variables (w) are winsorized at the 99th percentile. All monetary values are in 2021 USD in purchasing power parity terms. Heteroskedasticity-robust standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.8: Treatment Effects on Over- and Underconfident Beliefs - Both Experiments

	Tight experiment (quintiles)		Big experiment (terciles)	
	Underconfident (1)	Overconfident (2)	Underconfident (3)	Overconfident (4)
Treatment	-0.057*** (0.016)	-0.022 (0.018)	-0.043*** (0.005)	-0.101*** (0.007)
Treatment effect/ control mean	-0.306*** (0.087)	-0.035 (0.028)	-0.284*** (0.033)	-0.220*** (0.016)
p[Treat/mean(uc)]=p[Treat/mean(oc)]		0.001		0.000
Control mean	0.187	0.629	0.150	0.461
Observations	278	278	4205	4205

Notes: Table C.8 shows that underconfident beliefs are more likely to be updated than overconfident beliefs. It displays treatment effects on dummies indicating if the jobseeker's beliefs about her skills are underconfident (cols. 1 and 3) or overconfident (cols. 2 and 4). Cols. 1 to 2 show results for the tight experiment. Cols. 3 and 4 show results for the big experiment. The effect sizes divided by the control group means show that underconfident beliefs are more likely to update than underconfident beliefs. All specifications include randomization block fixed effects and prespecified baseline covariates described in footnotes 15 and 20. Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D Protocol and Intervention Details

This appendix summarizes the protocol and intervention details for the tight and the big experiment. Other relevant materials are in Appendix B (measurement of skills), and in Appendix C (measurement of beliefs about skills).

D.1 Tight Experiment

The data collection for the tight experiment ran from August to October 2022 in downtown Johannesburg. We recruited 373 participants using the user database of our implementation partner, the Harambee Youth Employment Accelerator. We contacted users who were active on our partner’s platform [SAYouth.mobi](#) in the past month, said they were residing in Gauteng province, had completed at least high school, and were at most 35 years old. Using this contact list, surveyors called potential participants, conducted a short screening survey, and invited them for a daylong job search workshop in the city center.³³ We offered 150 Rand (approximately 21 USD PPP) mobile phone airtime for compensation, which is equivalent to 6.5 hours work at the national minimum wage. The structure of data collection is depicted on Figure D.1 which we detail below.

When participants arrived at the venue, they received information about the schedule of the day, had breakfast, and were matched to a surveyor. The surveyor sought informed consent and started administering the surveys, programmed in Survey CTO, on a tablet. The surveyors were instructed to provide further explanations and translate the questions as needed, and to tailor the pace of the surveys to the needs of the participants.

The baseline survey collected participants’ demographic information, their employment and job search history, baseline beliefs about their skills and their labor market prospects, and their risk and time preferences. This survey was followed by three assessments: communication, numeracy, and concept formation, in that order. Participants had 30 minutes for each of the communication assessment and numeracy assessment and 15 minutes to complete the concept formation test. After the assessments, the surveyor administered a short survey about participants’ beliefs about their assessment performance.

On treatment days, the surveyors handed over the printed reports (Figure 1) to participants who then watched a video on the tablet with headphones. The video explained how participants should interpret the report, used several hypothetical examples for further explanation and prompted participants to review their own report and ask the surveyor any questions that they had. On control days, participants still viewed a minimally modified version of the video that omitted the explanation of the results. The scripts and

³³Prospective participants were screened out during this call if they were not currently searching for a job or were a full-time student.

the video were thoroughly piloted to ease participants' understanding. The treatment video is available at <https://bit.ly/3EoVoNL> and the control video is available at <https://bit.ly/3srwLgj>. The corresponding scripts are available at <https://bit.ly/45zthqu>. Baseline covariates are balanced across treatment arms (Table D.1).

After the treatment and a lunch break, the surveyors collected participants' beliefs about their skills and future labor market outcomes, and they administered the job choice task. In the job choice task, participants were asked to choose between two realistic jobs. Each job pair contained jobs with opposite skill demand (one communication- and one numeracy-heavy job) based on 13 recruiters' prior evaluation of the jobs. The jobs in the pairs were matched on several important dimensions: expected desirability, job-offer probability, and salary (as assessed by the recruiters), as well as location. These were all entry-level jobs that did not require certifications or equipment to ensure that all partic-

Table D.1: Balance Table - Tight Experiment

	Restricted sample					Full sample				
	Control (1)	Treatment (2)	Δ (3)	$p(\Delta = 0)$ (4)	N (5)	Control (6)	Treatment (7)	Δ (8)	$p(\Delta = 0)$ (9)	N (10)
<u>Panel A: Demographics</u>										
Black African	1.00	1.00	0.00	.	278	0.99	0.99	0.01	0.41	372
Male	0.33	0.32	-0.01	0.74	278	0.28	0.29	0.00	0.92	372
Age	26.42	26.40	-0.02	0.84	278	26.89	26.79	-0.09	0.91	372
Completed secondary education only	0.61	0.60	-0.01	0.68	278	0.62	0.58	-0.04	0.31	372
University degree / diploma	0.19	0.24	0.04	0.16	278	0.19	0.21	0.02	0.60	372
Any other post-secondary qualification	0.16	0.14	-0.01	0.46	278	0.15	0.17	0.02	0.42	372
<u>Panel B: Labor market background</u>										
Any work in last 7 days	0.35	0.31	-0.04	0.41	278	0.35	0.32	-0.03	0.49	372
Has worked in permanent wage job before	0.23	0.27	0.04	0.50	278	0.24	0.27	0.03	0.44	372
Earnings in USD (last 7 days, w)	46.28	43.53	-2.75	0.85	277	47.45	50.52	3.07	0.69	371
Written contract	0.09	0.16	0.06	0.04	278	0.11	0.18	0.07	0.03	372
<u>Panel C: Search behavior</u>										
Any job search in last 30 days	0.96	0.96	-0.01	0.72	278	0.96	0.96	-0.00	0.90	372
# applications (last 7 days, w)	11.31	8.69	-2.62	0.10	278	11.46	10.49	-0.97	0.54	372
Search expenditure in USD (last 7 days, w)	23.98	21.47	-2.51	0.25	278	24.42	22.65	-1.77	0.28	372
Hours spent searching (last 7 days, w)	13.75	13.90	0.15	0.92	278	14.06	14.23	0.17	0.88	371
# job offers (last 30 days, w)	0.14	0.21	0.07	0.11	278	0.20	0.23	0.03	0.51	372
<u>Panel D: Search alignment with CA</u>										
Δ planned apps (w, aligned - misaligned)	0.94	0.53	-0.41	0.55	278	0.57	0.40	-0.17	0.77	372
Δ % platform apps (aligned - misaligned)	-0.00	-0.00	0.00	0.73	278	-0.00	-0.00	0.00	0.74	372
<u>Panel E: Skills beliefs</u>										
Aligned belief about CA	0.47	0.50	0.04	0.68	278	0.43	0.46	0.03	0.69	369
Fraction aligned belief domains	0.18	0.26	0.08	0.12	278	0.18	0.24	0.06	0.17	369

Notes: **Table D.1** shows that covariates are balanced across treatment groups in the tight experiment. "CA" stands for comparative advantage in skills. Cols. 1 to 5 show results for the sample of individuals with a unique comparative advantage in skills. Cols. 6 to 10 show results for the full sample of individuals including those without unique comparative advantage in skills. P-values are estimated from regressions of each covariate on treatment and randomization block fixed effects using standard errors clustered by workshop day. Winsorized variables (w) are winsorized at the 99th percentile. All monetary values are in 2021 USD in purchasing power parity terms.

Table D.2: Job Titles Used in the Job Choice Task - Tight Experiment

Numeracy job title	Communication job title
Receiving and dispatching clerk	Sales agent
Sales teller	Customer service agent
Stock controller	General administrator
Laundry assistant	Waiter / waitress
Cashier	Host/hostess
Data capturer	Front desk assistant
Restaurant till manager	Receptionist
Store cashier	Sales assistant
Cash teller	Recruitment administrator
Banking call center agent	Retail call center agent
Petrol attendant	Maintenance assistant


Notes: **Table D.2** shows the job titles for numeracy and communication-heavy jobs in each job pair in the job choice task in the tight experiment. Each job pair contains one relatively numeracy-heavy and one communication-heavy job title. The pairs were created by matching jobs with different skill requirements but similar wages and general appeal to workers to ensure sufficient variation in job choices. The pairs are based on a list of 28 jobs rated by HR professionals for their skill requirements, expected wage, overall and gender-specific desirability, and transparency of skill requirements.

ipants could reasonably apply for them. See Table D.2 for the pairs of job titles. The job descriptions and their layout followed the design and style of the SAYouth.mobi platform. The job descriptions were presented side by side in a printed booklet to allow for easy comparisons. See Figure 2 for an example of one pair. Participants were shown the pairs in the booklet, read the descriptions and were asked to pick the job that they were most interested in applying to. Participants made decisions for the same set of 11 job pairs in a randomized order. The last two job pairs explicitly included the main skill (communication or numeracy) that the job required. For 5 of the 11 job pairs, after participants made their choice, we asked them the job offer probability and expected wage for each job if they were to apply to it.

We incentivized the job choices in two ways. First, one job pair of the 11 pairs included real job opportunities, and we submitted participants' application materials to the job that participants picked for this real pair. Participants were informed about this incentive, but they did not learn which pair was the real pair during the workshop. As a second incentive, participants received a list of job titles at the end of the workshop that matched their most preferred skill according to their choices to assist their job search.

In the next survey module, we measured participants' willingness-to-pay using incentivized multiple price lists for three products: for self-study materials to improve

Figure D.1: Tight Experiment Design



Notes: Figure D.1 shows the order of activities during each workshop in the tight experiment.

participants' communication skill and numeracy skill respectively, and a document that revealed the skill demand of a set of common jobs as per the rating of HR experts. Participants completed a practice round before the elicitation and the elicitation was simplified: we were asking a sequence of pairwise choice between a monetary value and the product itself. Further details of the protocol and the results are available at the following link: <https://bit.ly/3E34oYH>, which we have omitted from this appendix to save space.

After the willingness-to-pay module, we collected information typically found in CVs. Using this information we prepared a CV for each participant and delivered them to the participants after the workshop. We also submitted participants' CVs to the job that they preferred among the real jobs in the job choice task.

The final session of the workshop measured search effort. Participants had the opportunity to spend up to 15 minutes to write an optional short application email to the employer of the real jobs in the job choice task that they preferred. The message, along with participants' CV that we created, were both delivered to the employer, whose HR team could then evaluate the candidates. After this part, participants completed the check out procedure and received their compensation.

D.2 Big Experiment


The recruitment, data collection procedure and intervention in the big experiment were very similar to those in the tight experiment. In this section we highlight the main differences between the big and tight experiment. Figure D.2 shows the order of activities during workshops in the big experiment.

The big experiment took place in the same labor market, but earlier in time, in 2016/17. We again recruited participants from the contact list of the Harambee Youth Employment Accelerator who were aged 18–34, had completed secondary school, had at most 12 months of formal work experience, and were from disadvantaged backgrounds. The workshops took place at the Harambee office at the time, in downtown Johannesburg.

As in the tight experiment, jobseekers completed surveys and assessments. Treatment was randomized at the day level and the intervention consisted of receiving a report. The report (Figure D.3) showed the performance results for six skills (as opposed to only numeracy and communication in the tight experiment) and the results were reported in terciles and not quintiles. The report was handed to participants privately in an envelope and was followed by a briefing from industrial psychologists. The content of the briefing was very similar to that in the tight experiment video. The materials used for the briefing are available at <https://bit.ly/44o0p1v> and the script is available at <https://bit.ly/3YLaDK6>. Baseline covariates are balanced across treatment arms (Table D.3).

We conducted two follow-up surveys. A short SMS survey 2-3 days after the workshop, and a longer phone survey on average 3.5 months after the workshop. 96% of the sample was successfully interviewed at endline and we do not find differential attrition based on treatment (Table D.4). Respondents were compensated with mobile phone air-time for completing the phone survey.

Figure D.2: Big Experiment Design



Notes: Figure D.1 shows the order of activities during each workshop in the big experiment.

Figure D.3: Sample Treatment Report - Big Experiment

REPORT ON CANDIDATE COMPETENCIES

-Personal Copy-

This report contains results from the assessments you took at Harambee in Phase 1 and Phase 2. These results can help you learn about some of your strengths and weaknesses and inform your job search.

You completed assessments on English Communication (listening, reading and comprehension) and Numeracy today in Phase 2. In Phase 1, you completed a Concept Formation assessment which asked you to identify patterns.


1. The Numeracy tests measure various maths abilities. Your score is the average of the two maths tests you did today at Harambee.
2. The Communication test measures English language ability through listening, reading and comprehension.
3. The Concept Formation test measures the ability to understand and solve problems. Candidates with high scores can generally solve complex problems, while lower scores show an ability to solve less complex problems.

You also did some games and questionnaires to measure your soft skills:

4. The Planning Ability Test measures how you plan your actions in multi-step problems. Candidates with high scores generally plan one or more steps ahead in solving complex problems.
5. The Focus Test looks at your ability to pick out which information is important in confusing environments. Candidates with high scores are able to focus on tasks in distracting situations.
6. The Grit Scale measures candidates' determination when working on difficult problems. Candidates with high scores spend more time working on the problems rather than choosing to pursue different problems.

Your results have been compared to a large group of young South African job seekers who have a matric certificate, are from socially disadvantaged backgrounds and have been assessed by Harambee.

You scored in the MIDDLE THIRD of candidates assessed by Harambee for Numeracy, MIDDLE THIRD for Communication, LOWER THIRD for Concept Formation, LOWER THIRD for Planning Ability, MIDDLE THIRD for Focus and TOP THIRD for the Grit Scale.



DISCLAIMER

Please note that this is a confidential assessment report and is intended for use by the person specified above. Assessment results are not infallible and may not be entirely accurate.

Notes: Figure D.3 shows an example of the reports given to treated jobseekers in the big experiment. Each report contains the jobseeker's assessment results but no identifying information and no branding by Harambee.

Table D.3: Balance Table - Big Experiment

	Full sample					Non-attrited sample				
	Control (1)	Treatment (2)	Δ (3)	$p(\Delta = 0)$ (4)	N (5)	Control (6)	Treatment (7)	Δ (8)	$p(\Delta = 0)$ (9)	N (10)
Panel A: Demographics										
Black African	0.98	0.99	0.00	0.93	4389	0.98	0.99	0.00	0.98	4206
Male	0.39	0.36	-0.02	0.04	4389	0.39	0.36	-0.03	0.03	4206
Age	23.55	23.79	0.25	0.07	4389	23.53	23.80	0.26	0.05	4206
Completed secondary education only	0.62	0.59	-0.02	0.20	4389	0.62	0.59	-0.03	0.17	4206
University degree / diploma	0.16	0.18	0.02	0.29	4389	0.15	0.18	0.02	0.24	4206
Any other post-secondary qualification	0.21	0.22	0.01	0.61	4389	0.22	0.23	0.01	0.63	4206
Panel B: Labor market background										
Any worked in last 7 days	0.36	0.39	0.02	0.22	4389	0.36	0.39	0.03	0.19	4206
Has worked in permanent wage job before	0.09	0.09	-0.00	0.61	4377	0.10	0.09	-0.01	0.35	4195
Earnings in USD (last 7 days, w)	30.09	32.52	2.43	0.13	4389	29.84	32.49	2.65	0.11	4206
Panel C: Search behavior										
Any job search in last 7 days	0.97	0.97	0.01	0.07	4389	0.97	0.98	0.01	0.02	4206
# applications (last 30 days, w)	9.31	9.37	0.06	0.95	4346	9.13	9.27	0.14	0.79	4165
Search expenditure in USD (last 7 days, w)	32.09	31.14	-0.95	0.34	3912	32.02	31.15	-0.87	0.42	3747
Hours spent searching (last 7 days, w)	17.35	16.74	-0.61	0.32	4273	17.24	16.56	-0.68	0.28	4093
# job offers (last 30 days, w)	0.77	0.84	0.06	0.64	4335	0.78	0.87	0.09	0.44	4152
Panel D: Skills beliefs										
Aligned belief about CA	0.20	0.21	0.01	0.49	4312	0.20	0.21	0.01	0.67	4132
Fraction aligned belief domains	0.38	0.38	-0.00	0.49	4378	0.38	0.37	-0.00	0.45	4196

Notes: **Table D.3 shows that covariates are balanced across treatment arms for the big experiment.** “CA” stands for comparative advantage. Cols. 1 to 5 show statistics for the baseline sample. Cols. 6 to 10 show statistics for the 96% of participants reached at endline. P-values are estimated from regressions of each covariate on treatment and randomization block fixed effects using standard errors clustered by workshop day. Winsorized variables (w) are winsorized at the 99th percentile. All monetary values are in 2021 USD in purchasing power parity terms.

Table D.4: Treatment Effects on Completing Endline Survey - Big Experiment

	Interviewed at endline	
	(1)	(2)
Treatment	-0.005 (0.006)	-0.003 (0.006)
Control mean	0.960	0.960
Observations	4389	4389
Controls	No	Yes

Notes: **Table D.4 shows that attrition is low and balanced across treatment groups in the big experiment.** Dependent variable is a dummy indicating whether an individual was successfully recontacted for the endline survey. Column 2 includes prespecified baseline covariates described in footnote 20. Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E Additional Treatment Effects

E.1 Robustness Checks

This subsection shows robustness checks for changes in sample and variable definitions. Treatment effects on beliefs about skills (Table E.1) and search direction (Table E.2) are robust to including jobseekers without a unique comparative advantage, who are omitted in the main analysis of the tight experiment. Table E.3 shows the results of multiple robustness checks for key belief, search, and labor market outcomes in the tight experiment (cols. 1-2) and big experiment (cols. 3-6): estimating p-values using a wild cluster bootstrap, estimating sharpened q-values to account for multiple hypothesis testing, and reweighting the samples in the two experiments to have the same distribution of baseline covariates. Treatment effects on beliefs are robust to using non-binary measures of alignment between assessment results and beliefs about skills, both as outcomes and for heterogeneity analysis (Table E.4). Treatment effects on earnings are robust to different transformations (Table E.5) and are non-negative throughout the distribution (Figure E.1).

Table E.1: Treatment Effects on Skill Beliefs - Tight Experiment Including Jobseekers Without Unique Comparative Advantage

	Aligned CA belief		Fraction aligned beliefs	
	(1)	(2)	(3)	(4)
Treatment	0.120*** (0.037)	0.180*** (0.055)	0.089*** (0.020)	0.064*** (0.018)
Treatment $\times$ Aligned CA belief (bl)		-0.130 (0.089)		0.057 (0.041)
Aligned CA belief (bl)		0.447*** (0.081)		-0.039 (0.037)
Treatment effect: Aligned CA belief (bl)		0.050 (0.061)		0.121*** (0.036)
Control mean	0.446	0.446	0.171	0.171
Observations	368	368	368	368

Notes: Table E.1 shows that treatment effects on beliefs about skills are robust to including jobseekers who do not have a unique skill comparative advantage. “CA” stands for comparative advantage. Columns indicate different outcome variables: a dummy indicating whether the jobseeker’s beliefs about her comparative advantage in skills are aligned with the assessment results (cols. 1-2), and the fraction of her beliefs about her skill quintiles that are aligned with the assessment results (cols. 3-4). Cols. 2 and 4 show treatment effect heterogeneity by whether individuals had aligned comparative advantage beliefs at baseline. All regressions include prespecified control variables listed in footnote 15 and a dummy variable equal to one iff the jobseeker has no comparative advantage, which is prespecified but cannot be used when we restrict the sample to jobseekers with unique comparative advantage. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.2: Treatment Effects on Aligned Search Direction - Tight Experiment Including Jobseekers Without Unique Comparative Advantage

	% aligned (job choice)		Δ % aligned platform apps		Δ SMS click rate		Δ planned apps (w)		Aligned search index	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment	0.032 (0.025)	0.062** (0.026)	0.052*** (0.017)	0.062** (0.030)	0.071 (0.046)	0.122* (0.063)	1.121 (0.878)	3.363*** (1.140)	0.240*** (0.079)	0.436*** (0.094)
Treatment $\times$ Aligned CA belief (bl)		-0.070** (0.029)		-0.023 (0.060)		-0.113 (0.098)		-5.190** (2.162)		-0.446*** (0.156)
Aligned CA belief (bl)		0.111*** (0.026)		0.011 (0.035)		0.070 (0.076)		6.113*** (1.872)		0.506*** (0.121)
Treatment effect: Aligned CA belief (bl)		-0.008 (0.029)		0.040 (0.039)		0.009 (0.070)		-1.827 (1.613)		-0.011 (0.118)
Control mean	0.416	0.416	0.005	0.005	-0.024	-0.024	3.272	3.272	-0.000	-0.000
Observations	368	368	368	368	368	368	368	368	368	368

Notes: Table E.2 shows that treatment effects on aligned search direction are robust to including jobseekers who do not have a unique skill comparative advantage. All search direction measures are coded as zero for jobseekers with tied skill quintiles. “CA” stands for comparative advantage in skills. Aligned job search is defined as directing search effort toward jobs that mostly require the skill that aligns with jobseekers’ measured comparative advantage. Columns indicate different outcomes: the percentage of 11 incentivized job choices that are aligned with the measured comparative advantage of the jobseeker (cols. 1–2), the difference between the percentage of aligned and non-aligned applications on the online job search platform SAYouth.mobi (cols. 3–4), the difference in link click rates between aligned and non-aligned jobs sent to job seekers via SMS (cols. 5–6), the difference between aligned and non-aligned planned applications for the 30 days after the workshop (cols. 7–8), and the inverse covariance-weighted average of the measures displayed in cols. 1 to 8 following Anderson (2008) (cols. 9–10). Even-numbered columns show heterogeneity by whether individuals have aligned comparative advantage beliefs at baseline. All regressions include prespecified control variables listed in footnote 15 and a dummy variable equal to one iff the jobseeker has no comparative advantage, which is prespecified but cannot be used when we restrict the sample to jobseekers with unique comparative advantage. Winsorized variables (w) are winsorized at the 99th percentile. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.3: Robustness Checks for Main Outcomes - Comparing Tight and Big Experiments

	Tight experiment		Big experiment			
	Aligned CA belief (1)	Aligned search index (2)	Aligned CA belief (3)	Aligned search direction (4)	LM quantity index (5)	LM quality index (6)
<u>Panel A: Robustness for inference</u>						
Treatment	0.135	0.254	0.139	0.050	0.045	0.085
clustered p-value	(0.001)	(0.017)	(0.000)	(0.000)	(0.165)	(0.017)
wild-bootstrapped p-value	[0.005]	[0.009]	[0.000]	[0.000]	[0.395]	[0.005]
q-value	{0.002}	{0.011}	{0.001}	{0.001}	{0.029}	{0.011}
<u>Panel B: Reweighting to match other sample</u>						
Treatment	0.149*** (0.044)	0.254** (0.101)	0.144*** (0.011)	0.054*** (0.010)	0.056* (0.032)	0.086** (0.037)
Control mean	0.475	0.000	0.196	0.165	-0.000	0.000
Observations	278	278	4118	4205	4205	4206
Number of clusters	34	34	54	54	54	54

Notes: **Table E.3** shows that the main results are robust to different hypothesis testing approaches and to reweighting data to match the demographics across experiments. **Panel A** shows p-values from cluster-robust standard errors, p-values from a wild cluster bootstrap with 10,000 replications, and sharpened q-values that control the false discovery rate across tests on all outcomes shown in the table, following [Benjamini et al. \(2006\)](#). **Panel B** shows the same treatment effects estimated with the experimental samples reweighted to have the same distribution of baseline covariates. The covariates used to estimate weights are all variables displayed in the summary statistics [Tables 1](#) and [A.2](#) that are consistently measured in both experiments: dummies for being black, male, having only completed secondary education, having a post-secondary degree, or having a post-secondary certificate, having worked in the last seven days, and having ever worked in a wage job, as well as continuous measures of earnings, the number of job applications, job search hours and expenditure, and job offers in the last seven days. Cols. 1–2 show effects on the main outcomes of the tight experiment. Cols. 3–6 show effects on the main outcomes of the big experiment. Control variables for the tight experiment are listed in footnote [15](#). Control variables for the big experiment are listed in footnote [20](#). Standard errors clustered at the treatment-day level in parentheses in panel B. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ in Panel B.

Table E.4: Heterogeneous Treatment Effects by Nonbinary Baseline Comparative Advantage Beliefs - Tight Experiment

	Main outcomes				Additional belief outcomes			
	Aligned search index		Aligned CA belief		Degree of CA alignment		Overall alignment of belief levels	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	0.269** (0.103)	0.842* (0.414)	0.152*** (0.044)	0.444*** (0.142)	0.063*** (0.012)	0.324*** (0.091)	0.073*** (0.012)	0.138** (0.062)
Treatment $\times$ Baseline degree of CA belief alignment		-0.681 (0.504)		-0.348* (0.176)		-0.315*** (0.106)		-0.078 (0.072)
Baseline degree of CA belief alignment		1.595*** (0.495)		1.125*** (0.215)		0.828*** (0.071)		0.045 (0.050)
Control mean	-0.000	-0.000	0.475	0.475	0.816	0.816	0.511	0.511
Observations	278	278	278	278	278	278	278	278


Notes: Table E.4 shows that average and heterogeneous treatment effects are robust to using nonbinary measures of aligned skill beliefs in the tight experiment. The outcomes in cols 1–4 are the same as those used in the main analysis. “Degree of CA belief alignment” captures the degree to which comparative advantage beliefs are aligned. It is defined as $1 - \text{abs}(\Delta \text{bel_skill}_{i,j} + 1) / 5$, where $\Delta \text{bel_skill}_{i,j}$ is the difference between numeracy and communication beliefs (in quintiles). Hence $\text{abs}(\Delta \text{bel_skill}_{i,j} + 1)$ shows the minimum number of quintiles the jobseeker needs to adjust their skill beliefs to align them with their measured comparative advantage beliefs. It is equal to 0 when the jobseeker’s comparative advantage belief is aligned (as $\Delta \text{bel_skill}_{i,j} = 0$) and has a maximum value of 1. This is the outcome in cols. 5–6 and is used for heterogeneity analysis in all even-numbered columns. The mean of this measure is 0.83 and the standard deviation is 0.22. “Overall alignment of belief levels” is defined as one minus the average difference between beliefs about skill quintiles and assessed quintiles divided by the maximum possible difference. A value of 1 indicates perfectly aligned beliefs while a value of 0 indicates maximally misaligned beliefs. This is the outcome in cols 7–8. Control variables are listed in footnote 15. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.5: Treatment Effects on Different Transformations of Earnings - Big Experiment

	Unconditional earnings					Conditional earnings				
	raw (1)	wins (2)	z (3)	ln(earn+1) (4)	ih (5)	raw (6)	wins (7)	z (8)	ln(earn+1) (9)	ih (10)
Treatment	9.360** (3.610)	6.517** (2.712)	0.097** (0.037)	0.115** (0.053)	0.127** (0.060)	25.100** (11.348)	20.393** (9.404)	0.159** (0.072)	0.229** (0.093)	0.244** (0.101)
Control mean	27.080	25.424	-0.000	0.954	1.109	88.109	85.826	-0.000	3.105	3.608
Observations	4196	4196	4196	4196	4196	1280	1280	1280	1280	1280

Notes: **Table E.5 shows that treatment effects on earnings in the seven days before the endline survey are robust to a range of transformations.** Columns show effects on transformed unconditional earnings variable (cols. 1-5) and on transformed earnings variable conditional on doing any work (cols. 6-10). Outcomes in each column are raw earnings (cols. 1 and 6), earnings winsorized earnings at the 99th percentile (cols. 2 and 7), earnings standardized with respect to the control group mean and standard deviation (cols. 3 and 8), the natural logarithm of earnings plus one (cols. 4 and 9), and the inverse hyperbolic sine transformation of earnings (cols. 5 and 10). Control variables are listed in footnote 20. Earnings are measured in 2021 USD purchasing power parity terms. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure E.1: Cumulative Distributions of Earnings by Treatment Group - Big Experiment



Notes: **Figure E.1 shows that treatment effects on earnings in the big experiment reflect a rightward shift in the entire distribution.** It shows the cumulative distribution function of earnings in the last seven days in the big experiment for the treatment and control groups. Earnings are measured in 2021 USD purchasing power parity terms and are winsorized at the 99th percentile. Panel A shows earnings conditional on having worked in the last seven days. Panel B shows unconditional earnings for the full sample.

E.2 Heterogenous Treatment Effects

This section shows heterogeneous treatment effects mentioned in the text of the main paper. Table E.6 shows heterogeneous treatment effects on beliefs about skills and search direction by baseline aligned comparative advantage belief in the big experiment. These are consistent with the results of tight experiment. Table E.7 shows heterogeneous treatment effects on beliefs about skills, search direction, and search effort by baseline aligned comparative advantage beliefs and baseline confidence in the tight experiment. These show that the heterogeneous treatment effects by baseline aligned comparative advantage beliefs are broadly robust to allowing effects to vary also by baseline confidence. Table E.8 shows heterogeneous treatment effects on search direction in the tight experiment's job choice task by whether skill requirements for specific job pairs were revealed. This shows that effects on search direction are larger when skill requirements are revealed but are positive even when they are not revealed. Table E.9 show treatment effects on job quality outcomes (earnings and written contract) multiplied by potential mediators, such as tenure and start date. Finally, Table E.10 shows heterogeneous treatment effects on labor market outcomes by aligned baseline comparative advantage beliefs in the big experiment. These are also consistent with the results of the tight experiment.

Table E.6: Heterogeneous Treatment Effects on Beliefs about Skills and Search Direction by Aligned Baseline Comparative Advantage Belief - Big Experiment

	Beliefs		Search direction
	Aligned CA belief (1)	Fraction aligned beliefs (2)	Aligned search (3)
Treatment	0.142*** (0.011)	0.134*** (0.009)	0.060*** (0.012)
Treatment $\times$ Aligned CA belief (bl)	-0.014 (0.036)	0.043** (0.020)	-0.057* (0.032)
Aligned CA belief (bl)	0.157*** (0.027)	-0.008 (0.013)	-0.013 (0.024)
Treatment effect: Aligned CA belief (bl)	0.129*** (0.034)	0.177*** (0.020)	0.003 (0.028)
Control mean	0.196	0.388	0.165
Observations	4118	4131	4131

Notes: Table E.6 shows that treatment effects by baseline aligned comparative advantage beliefs on skill beliefs and search direction in the big experiment are broadly consistent with the results of the tight experiment. "CA" stands for comparative advantage in skills and "bl" stands for baseline. This table uses the same three outcomes as Table 4: a dummy for aligned comparative advantage beliefs (col. 1), the fraction of aligned belief domains (col. 2), and a dummy indicating if self-reported search direction is aligned with comparative advantage in skills (col. 3). Control variables are listed in footnote 20. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.7: Heterogeneous Treatment Effects on Main Outcomes by Baseline Aligned Comparative Advantage Beliefs and Confidence - Tight Experiment

	Aligned CA belief (1)	Aligned search index (2)	Search effort index (3)
Treatment	0.216*** (0.066)	0.550*** (0.155)	-0.192 (0.177)
Treatment $\times$ Aligned CA belief (bl)	-0.161 (0.112)	-0.651** (0.247)	0.274 (0.298)
Aligned CA belief (bl)	0.737*** (0.110)	0.539** (0.211)	-0.053 (0.226)
Treatment $\times$ Average confidence (bl)	-0.003 (0.079)	0.084 (0.179)	0.081 (0.175)
Average confidence (bl)	0.072 (0.058)	0.061 (0.142)	0.201 (0.143)
Treatment $\times$ Aligned CA belief (bl) $\times$ Average confidence (bl)	0.013 (0.111)	-0.065 (0.208)	-0.167 (0.230)
Control mean	0.475	0.000	0.000
Observations	278	278	278

Notes: Table E.7 shows that the main results are robust to interacting the treatment with both baseline aligned comparative advantage beliefs and baseline confidence in skills. The magnitudes of the interaction terms in row 2 are very similar to the estimates from regressions that exclude baseline confidence, shown in Tables 2 and 3. The estimated coefficients on the three-way interaction term (row 6) are small and not statistically significant. “CA” stands for comparative advantage in skills and “bl” stands for baseline. Baseline confidence levels are defined as the average baseline deviation of beliefs about skill quintiles from measured quintiles (across numeracy and communication), divided by the control group standard deviation. Positive values indicate overconfidence and negative values indicate underconfidence. Columns indicate different outcome variables: a dummy indicating aligned comparative advantage beliefs (col. 1), the aligned search index (col. 2), and the search effort index (col. 3). Control variables are listed in footnote 15. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.8: Heterogeneous Treatment Effects on Search Direction for Jobs With and Without Revealed Skill Demands - Tight Experiment

	Choice aligned with measured comp. adv.	
	(1)	(2)
Treatment	0.039 (0.032)	0.023 (0.034)
Treatment $\times$ Skill req. revealed	0.130** (0.063)	-0.114 (0.085)
Treatment effect: Skill req. revealed	0.169*** (0.060)	-0.091 (0.077)
Control mean	0.550	0.550
Observations	1573	1485
Sample: jobseekers with baseline CA beliefs	misaligned	aligned

Notes: Table E.8 shows that, for jobseekers with misaligned baseline comparative advantage beliefs, treatment effects on search direction in the job choice task are stronger when jobs' skill requirements are revealed. The outcome variable of the regressions is a dummy variable indicating if job choices in the job choice task are aligned with jobseekers' assessed comparative advantage. The effects are estimated at the job-pair $\times$ individual level for jobseekers with misaligned baseline beliefs about comparative advantage (col. 1) and for jobseekers with aligned baseline beliefs (col. 2). Controls include randomization block fixed effects, job pair and job pair order fixed effects, and prespecified baseline covariates described in footnote 15. "CA" stands for comparative advantage in skills. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.9: Treatment Effects on Combinations of Labor Market Outcome Measures - Big Experiment

	Earnings in the last seven days (w)					Written contract	
	Start before treatment (1)	Start after treatment (2)	Tenure (3)	Wage emp. (4)	Self emp. (5)	Start before treatment (6)	Start after treatment (7)
Treatment	1.985* (1.072)	4.607** (2.237)	10.303 (8.297)	6.923*** (2.562)	-0.405 (0.694)	0.003 (0.004)	0.014 (0.010)
Control mean	4.164	21.095	51.369	18.463	4.184	0.018	0.102
Observations	4176	4176	4173	4174	4174	4184	4184

Notes: Table E.9 shows that treatment effects on labor market outcomes are driven by wage jobs started after the treatment. It displays treatment effects on selected work quality outcomes (earnings and written contract) multiplied by potential mediators (row 2 headers). Control variables are listed in footnote 20. Winsorized variables (w) are winsorized at the 99th percentile. All monetary values are measured in 2021 USD in purchasing power parity terms. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.10: Heterogeneous Treatment Effects on Labor Market Outcomes by Aligned Baseline Comparative Advantage Beliefs - Big Experiment

	Work quantity					Work quality			
	# job offers (w) (1)	Worked month 1 (2)	Worked month 2 (3)	Worked last 7 days (4)	Index (5)	Earnings (w) (6)	Hourly wage (w) (7)	Written contract (8)	Index (9)
Treatment	0.036 (0.024)	0.037** (0.015)	0.009 (0.018)	0.013 (0.014)	0.065* (0.035)	7.467** (3.230)	0.355** (0.149)	0.018 (0.011)	0.097** (0.041)
Treatment $\times$ Aligned CA belief (bl)	-0.023 (0.044)	-0.057 (0.042)	-0.015 (0.034)	-0.017 (0.031)	-0.079 (0.075)	-4.697 (5.261)	-0.304 (0.251)	-0.002 (0.022)	-0.054 (0.062)
Aligned CA belief (bl)	0.009 (0.032)	0.045 (0.028)	0.018 (0.024)	-0.014 (0.019)	0.053 (0.055)	2.796 (3.577)	0.244 (0.178)	-0.023 (0.015)	-0.004 (0.044)
Treatment effect: Aligned CA belief (bl)	0.013 (0.043)	-0.019 (0.038)	-0.006 (0.029)	-0.004 (0.027)	-0.015 (0.069)	2.770 (4.140)	0.051 (0.219)	0.016 (0.018)	0.043 (0.048)
Control mean	0.195	0.465	0.437	0.309	-0.000	25.424	1.267	0.120	0.000
Observations	4071	4127	4130	4130	4131	4122	4111	4111	4132

Notes: Table E.10 shows that treatment effects on labor market outcomes in the big experiment are driven by jobseekers with misaligned baseline comparative advantage beliefs. “CA” stands for comparative advantage in skills and “bl” stands for baseline. Outcomes are the same as in Table 5: the winsorized number of job offers in the preceding 30 days (col. 1), a dummy indicating any work for pay 1 month after treatment (col. 2), a dummy indicating any work for pay 2 months after treatment (col. 3), a dummy indicating any work for pay in the preceding seven days (col. 4), an inverse covariance-weighted average index of the four employment quantity measures shown in cols. 1–4 (col. 5), winsorized earnings in the preceding seven days (col. 6), winsorized hourly wages in the preceding seven days (col. 7), a dummy indicating a written contract (col. 8), and an inverse covariance-weighted average of the three employment quality measures shown in cols. 6–8 (col. 9). Control variables are listed in footnote 20. Winsorized variables (w) are winsorized at the 99th percentile. All monetary values are measured in 2021 USD in purchasing power parity terms. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F Beliefs About Wages and Job Offer Probabilities

The model predicts a three-stage reaction to treatment: beliefs about skill comparative advantage → search direction → search outcomes. We show evidence consistent with each of these stages in the paper. The model also suggests that treatment will shift expected returns to specific search activities. We evaluate this idea in this appendix.

We find that treatment effects on beliefs about search outcomes are broadly consistent with the model. But we view them as less important than the results on beliefs about skills, search direction, and search outcomes we include in the main paper. The questions about expected search outcomes rely on complex forecasts by jobseekers, as discussed in recent reviews ([Delavande, 2023](#); [Mueller & Spinnewijn, 2023](#)). For example, we ask jobseekers about their expected search duration. This requires jobseekers to forecast both the expected number of offers and attributes of those offers (wages, hours, travel costs, working conditions, etc.), as the attributes determine whether they will accept offers. In contrast, our tight experiment is optimized to measure search direction in multiple different ways, including direct measures of behavior. Moreover, the survey questions about skill beliefs and actual labor market outcomes are simpler to answer.

We use only data from the tight experiment in this section. In the tight experiment, we collected beliefs about search outcomes on the same day as treatment. This means that treatment effects on beliefs reflect only information acquired during the workshops. In the big experiment, the endline survey occurred months after treatment. This means that any treatment effects on beliefs about search outcomes would reflect both the direct effect of treatment and any indirect effects arising from treatment-induced changes in search actions and their outcomes. The indirect causal channel is interesting but not a key part of the argument that we make in this paper.

Treatment effects on beliefs about returns to search direction: The model suggests that changes in search direction will be accompanied by changes in beliefs about the relative returns to searching for different types of jobs, i.e., beliefs about the returns to skill-directed job search. To evaluate this prediction, we conduct two exercises.

First, we surveyed jobseekers about their expected outcomes from applying to each type of job and estimate treatment effects on these measures. We asked for their expected number of job offers in the next 30 days, time to employment, and wage when employed. We asked these questions after treatment and just after asking their planned number of applications in the next 30 days, and we ask the expected offers and time to employment questions conditional on their planned number of applications. We asked all questions separately about all jobs, communication-heavy jobs, and numeracy-heavy

Table F.1: Treatment Effects on Beliefs About Returns to Skill-Directed Job Search - Tight Experiment

	Index		ΔE [offers per app] (w)		$-(\Delta [\text{sear. dur.}])$ (w)		ΔE [wage] (w)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	0.095 (0.100)	0.260* (0.130)	0.043** (0.018)	0.053** (0.026)	-0.007 (0.157)	-0.079 (0.185)	0.073 (6.605)	23.981** (9.383)
Treatment $\times$ Aligned CA belief (bl)		-0.341* (0.180)		-0.023 (0.034)		0.114 (0.223)		-45.738*** (13.960)
Aligned CA belief (bl)		0.689*** (0.140)		0.107*** (0.030)		0.432** (0.197)		35.845*** (12.069)
Treatment effect: Aligned CA belief (bl)		-0.081 (0.124)		0.030 (0.022)		0.035 (0.191)		-21.757** (9.164)
Control mean	-0.000	-0.000	0.027	0.027	0.244	0.244	6.007	6.007
Observations	278	278	273	273	272	272	278	278

Notes: Table F.1 shows that treatment improves jobseekers' their expected outcomes from searching for jobs that match their assessed comparative advantage relative to jobs that do not match their assessed comparative advantage. All outcomes are defined as the expected outcome for jobs aligned with comparative advantage minus the expected outcome for jobs not aligned with comparative advantage, so positive values indicate higher expected returns from searching for aligned jobs. "CA" stands for comparative advantage in skills and "bl" stands for baseline. Outcomes are an inverse covariance-weighted average of outcomes in cols 3–8 (cols. 1–2) and differences, as defined above, in winsorized expected offers per application (cols. 3–4), winsorized expected search duration in months (cols. 5–6), and winsorized expected weekly wages in 2021 USD in purchasing power parity terms (cols. 7–8). Even-numbered columns show heterogeneity by whether individuals have aligned comparative advantage beliefs at baseline. Control variables are listed in footnote 15. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

jobs. We define the expected return to skill-directed job search as the expected outcome for jobs aligned with comparative advantage minus the expected outcome for jobs not aligned with comparative advantage. For example, for a jobseeker with numeracy comparative advantage, the expected number of offers for numeracy-heavy jobs minus the expected number of offers for communication-heavy jobs. We also construct an inverse covariance-weighted average of the three measures.

Treatment effects on these survey measures of expected returns to skill-directed job search are mostly consistent with our model. Treated jobseekers expect 0.043 more job offers from skill-directed job search, equal to 165% of the control group mean (Table F.1, column 3, row 1, $p = 0.023$). This effect, like most in the paper, is driven by jobseekers with baseline misaligned comparative advantage beliefs (column 4). Treated jobseekers with baseline misaligned comparative advantage beliefs expect weekly earnings 24 USD higher from skill-directed job search, equal to roughly four times the control group mean (column 8, row 1, $p < 0.001$). Treatment has negligible effects on the expected length of time it will take to get a job from skill-directed job search (columns 5–6). We return to this

result on page 81.

Treatment effects on an index combining these survey measures of expected returns to skill-directed job search are consistent with our model, although slightly imprecisely estimated. Treatment increases the expected return to skill-directed search by 0.095 standard deviations for the average jobseeker (column 1, row 1, $p = 0.35$). This result is driven by the same heterogeneity that we see elsewhere in the paper: jobseekers with misaligned baseline comparative advantage beliefs increase their expected returns by 0.26 standard deviations (column 2, row 1, $p = 0.054$) while jobseekers with aligned baseline comparative advantage beliefs do not increase their expected returns (column 2, row 4).

Reassuringly, we see a “sensible” relationship between baseline comparative advantage beliefs and expected returns to skill-directed job search in the control group. Jobseekers whose baseline comparative advantage belief matches their assessed comparative advantage have a 0.69 standard deviation higher expected return to skill-directed job search (column 2, row 3). The relationship is also positive for all three components of the index (columns 4, 6, 8).

Second, we survey jobseekers about their expected outcomes from applying to specific jobs during the job choice task. For each job in 5 of the 11 job pairs, we ask jobseekers about the probability of getting an offer if they applied, the expected starting wage, and the general desirability of the job. We estimate treatment effects on these measures with a prespecified regression of the belief measure on treatment, a dummy for job alignment with jobseeker comparative advantage, their interaction, job fixed effects, and prespecified controls. The relevant coefficient is on the interaction term. This captures the treatment effect on jobseekers’ beliefs about the attributes of the job aligned with their skill comparative advantage, relative to the job not aligned with their comparative advantage.

Treatment effects on these measures of expectations are similar to effects on the first types of expectations, discussed above, but are less precisely estimated. Within each pair of jobs, treatment increases the expected offer probability and wage for the job aligned with the jobseeker’s comparative advantage (Table F.2, columns 2–3). But the effects are small – roughly 2% of the control group mean – and are not statistically significant at conventional levels. So we do not view this as strong evidence supporting the model. These results might be less precise than the results using the survey measures of expected returns to skill-directed search discussed above because the questions asked during the job choice task only ask about five specific pairs of jobs, rather than allowing jobseekers to implicitly average over many skill-directed job application choices.

Relationship between search direction and expected return to skill-directed job search: Beliefs about job attributes predict jobseekers’ choices in the job choice task, con-

Table F.2: Treatment Effects on Beliefs About Jobs in the Job Choice Task - Tight Experiment

	Desirability (sd)	Expected earnings (w)	Job offer probability
	(1)	(2)	(3)
Treatment	-0.038 (0.032)	-3.240 (5.921)	-0.022 (0.022)
Treatment $\times$ Aligned skill req.	0.008 (0.029)	3.873 (4.006)	0.011 (0.013)
Aligned skill req.	0.030 (0.022)	-3.297 (2.724)	0.017 (0.011)
Control mean	-0.000	193.588	0.544
Observations	2770	2770	2770

Notes: Table F.2 shows that treatment very weakly increased jobseekers' expectations about jobs in the job choice task that were aligned with their assessed comparative advantage. Beliefs were elicited for both jobs in each of 5 job pairs for each jobseeker. Columns indicate different outcome variables: desirability of jobs measured on a 0–10 Likert scale and standardized with respect to the control group (col. 1), expected weekly earnings winsorized at the 99th percentile (col. 2), and expected job offer probability (col. 3). Analysis is at the job $\times$ jobseeker level and includes prespecified control variables listed in footnote 15, dummies for comparative advantage in numeracy and communication, fixed effects for the randomized order in which job pairs were shown, and job fixed effects. All monetary values are measured in 2021 USD in purchasing power parity terms. Standard errors clustered at the treatment-day level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

sistent with a role for belief-based job search. For 5 of the 11 pairs of jobs, we asked jobseekers the probability that they would get an offer if they applied to each job and their expected salary if offered a job. We regress the job choice on the job offer probability times the expected monthly wage using a logit regression model, following Wiswall & Zafar (2015). This is not an experimental analysis because we regress post-treatment choices on post-treatment beliefs. Column 1 of Table F.3 shows that a 100 USD increase in the expected weekly wage scaled by the offer probability is associated with a 9.2 pp increase in the probability of choosing that job ($p = 0.001$). This relationship is robust to adding both jobseeker and job pair fixed effects (columns 2–4). The slope of the belief-choice relationship is somewhat different in the treatment and control groups, although the differences are not statistically significant. This suggests that treatment effects on these belief measures account for part but not all of the treatment effects on job choices.

Treatment effects on beliefs about search outcomes: The preceding analysis focuses on jobseekers' beliefs about the returns to searching for jobs requiring specific skills. We also construct measures of jobseekers' expected outcomes from searching for any type of job. We estimate treatment effects on these beliefs using the specification in equation (4).

Treatment has a positive but imprecisely estimated effect on expected wages, driven

by jobseekers with baseline misaligned comparative advantage beliefs. Columns 3–4 of Table F.4 show that treatment increases expected weekly wage by 9.6 USD for the average treated jobseeker and 22.1 USD for the average treated jobseeker with misaligned baseline comparative advantage belief. Both effects are relatively imprecisely estimated (standard errors 7.7 and 14.9 respectively), perhaps reflecting the difficulty of forecasting wage offers for respondents with limited work experience. Effects on jobseekers’ reservation wages, beliefs about the minimum and maximum wages they might earn if employed, and an index combining all of these wage beliefs follow the same qualitative pattern (columns 1–2 and 5–10).³⁴ These positive effects on wage beliefs in the tight experiment are consistent with the positive effects on actual wages in the big experiment, although we do not compare the magnitudes to evaluate forecast accuracy because the estimates come from two different experiments.

Treatment has a positive effect on the expected probability of formal employment, again driven by jobseekers with baseline misaligned comparative advantage beliefs. Columns 9–12 of Table F.5 show that treatment increases the probability of employment in 1–3 months by 3–4 percentage points for the average treated jobseeker and 8–9 percentage points for the average treated jobseeker with misaligned baseline comparative advantage belief. Effects on other, less direct proxies for employment probability – callbacks and offers per application and search duration – are closer to zero (columns 3–8). These results might differ because the callback, offer, and search duration questions all explicitly condition on the jobseeker’s planned number of applications, while the probability of employment questions are asked later in the survey and do not include this explicit conditioning.³⁵ The explicit conditioning might mean jobseekers put more mental weight on the role of search effort relative to search direction when answering these questions, but this is a speculative suggestion that future work could better evaluate. These results from the tight experiment are qualitatively consistent with the big experiment’s positive effect on employment with a written contract. However, the magnitudes are substantially different, perhaps in part because the control group’s expectations are much higher than realized outcomes.

³⁴In standard job search models, the reservation wage is both a decision rule and a feature of the wage distribution. This is not inconsistent with our interpretation of reservation wages as another proxy for wage expectations.

³⁵For example, we ask “How many job applications do you plan to submit in the next 30 days?” and then “If you submit X job applications in the next 30 days, how many months starting from today do you think it will take you to find a formal job, with an employment contract where you are paid a regular salary?”

Table F.3: Association between Job Choices, Offer Probabilities, and Expected Wages - Tight Experiment

	Marginal effects on choice of numeracy job (logit estimate)							
	Control group				Control and treatment groups			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$E[Wage^{num}] - E[Wage^{com}] (w)$	0.092*** (0.028)	0.097*** (0.030)	0.119*** (0.038)	0.127*** (0.040)	0.092** (0.038)	0.099** (0.039)	0.120*** (0.038)	0.124*** (0.039)
Treatment $\times E[Wage^{num}] - E[Wage^{com}] (w)$					0.023 (0.052)	0.022 (0.051)	-0.076 (0.055)	-0.078 (0.055)
Observations	695	695	600	600	1385	1385	1195	1195
Job pair fixed effects	No	Yes	No	Yes	No	Yes	No	Yes
Individual fixed effects	No	No	Yes	Yes	No	No	Yes	Yes

Notes: **Table F.3 shows that expected earnings predict job choice in the tight experiment.** All estimates are average marginal effects from logit regressions of indicators for choosing the numeracy-heavy job on the difference in expected returns between the communication-heavy and numeracy-heavy jobs. The difference in expected returns is the expected weekly wage (in 100s of USD in 2021 purchasing power parity terms) times the expected probability of a job offer for the numeracy job, minus the equivalent quantity for the communication job, winsorized at the 1% and 99% levels. Columns 1–4 use only the control group and columns 5–8 use both the treatment and control groups. Columns 3–4 and 7–8 include jobseeker fixed effects. Sample sizes drop from columns 1–2 and 5–6 to columns 3–4 and 7–8 because some jobseekers choose the numeracy-heavy job in all pairs, so their fixed effects are perfect predictors. Standard errors shown in parentheses are heteroskedasticity-robust in all columns, clustering by treatment date where appropriate, and using a 500-repetition bootstrap in columns with fixed effects, where analytical heteroskedasticity-robust adjustments are not feasible. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F.4: Treatment Effects on Beliefs about Wages - Tight Experiment

	Index		Wage expectations (w)		Minimum expected wage (w)		Maximum expected wage (w)		Reservation wage (w)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment	0.099 (0.104)	0.319* (0.179)	9.566 (7.740)	22.119 (14.852)	8.934 (6.090)	19.180* (10.607)	0.650 (11.984)	3.115 (24.259)	1.318 (3.628)	3.805 (5.948)
Treatment $\times$ Aligned CA belief (bl)		-0.461* (0.258)		-26.041 (21.034)		-21.065 (17.531)		-6.425 (38.832)		-5.366 (6.578)
Aligned CA belief (bl)		0.247 (0.208)		10.018 (17.725)		4.604 (15.279)		20.878 (29.176)		5.287 (4.039)
Control mean	-0.000	-0.000	212.045	212.045	131.392	131.392	324.076	324.076	109.567	109.567
Observations	278	278	278	278	278	278	278	278	277	277

Notes: Table F.4 shows that treatment effects on beliefs about wages in the tight experiment are positive but not consistently statistically significant. “CA” stands for comparative advantage in skills and “bl” stands for baseline. Columns indicate different outcome variables: an inverse covariance-weighted average of the other outcomes (cols. 1–2), the expected wage (cols. 3–4), the minimum expected wage (cols. 5–6), the maximum expected wage (cols. 7–8), and the reservation wage (cols. 9–10). Wage expectations are answered to the questions “What is the (lowest / highest possible) monthly take-home salary you think you would earn?” in a permanent, formal job with six months tenure. Even-numbered columns show heterogeneity by whether jobseekers had aligned comparative advantage beliefs at baseline. Control variables are listed in footnote 15. All outcomes are winsorized at the 99th percentile. All monetary values are measured in 2021 USD in purchasing power parity terms per week. Standard errors clustered at the treatment-day level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F.5: Treatment Effects on Beliefs about Offers and Search Duration - Tight Experiment

	Index		Callbacks / apps (w)		Offers / apps (w)		Month to job (w)		p(employed in 1 months)		p(employed in 3 months)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treatment	0.095 (0.122)	0.264 (0.194)	-0.001 (0.023)	0.039 (0.036)	-0.006 (0.021)	0.016 (0.031)	0.011 (0.247)	0.228 (0.392)	3.493 (2.238)	8.674** (3.748)	4.235* (2.315)	8.431** (4.042)
Treatment $\times$ Aligned CA belief (bl)		-0.360 (0.247)		-0.086 (0.057)		-0.049 (0.042)		-0.444 (0.470)		-10.821* (6.003)		-8.678* (5.126)
Aligned CA belief (bl)		0.234 (0.228)		0.061 (0.043)		0.049 (0.036)		0.069 (0.386)		4.789 (4.416)		2.647 (4.186)
Control mean	0.000	0.000	0.391	0.391	0.257	0.257	2.466	2.466	54.094	54.094	68.230	68.230
Observations	278	278	276	276	274	274	275	275	278	278	278	278

Notes: Table F.5 shows that treatment effects on beliefs about job offers and search duration in the tight experiment are mostly positive but not consistently statistically significant. “CA” stands for comparative advantage in skills and “bl” stands for baseline. Columns indicate different outcome variables: an inverse covariance-weighted average of the other outcomes (cols. 1–2), expected callbacks per application (cols. 3–4), expected offers per application (cols. 5–6), expected months to find a full-time job (cols. 7–8), and expected probability of being employed one month (cols. 9–10) and three months after baseline (cols. 11–12). Expected months to employment is multiplied by minus one before being included in the average so that higher values correspond to better search outcomes, as for the other outcomes. Winsorized variables (w) are winsorized at the 99th percentile. Even-numbered columns show heterogeneity by whether jobseekers had aligned comparative advantage beliefs at baseline. Control variables are listed in footnote 15. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

G Benefit-Cost Comparison

This appendix reports the variable costs of the assessment operation and compares these to the earnings gains experienced by treated jobseekers in the big experiment, showing that the gains likely exceed the costs.

The average treatment effect on weekly earnings was 6.52 USD in 2021 purchasing power parity terms at the time of the follow-up survey, which occurred an average of 14.5 weeks after treatment. Using this to forecast the average lifetime earnings gain for treated jobseekers requires very strong modeling assumptions. Instead, we take a conservative approach and assume that the treatment effect is constant from treatment to endline and zero thereafter. This implies an average treatment effect of 94.50 USD on earnings.

We calculate that the average variable cost of the assessment operation is 45.74 USD using data from Harambee and J-PAL Africa's accounting records. This consists of 12.17 for rent and utilities for the assessment center; 10.14 for depreciation of the computers used in assessment; 0.35 for assessment and software licenses; 8.11 for airtime and data, mainly for internet access for the assessment computers and for contacting jobseekers; 3.82 for Harambee salaries for assessment support staff, psychologists delivering briefings about results, and administrative support; 1.40 for J-PAL Africa salaries for research staff who helped to run the assessments and results briefings; and 9.74 for transport money for jobseekers to attend the assessments.

This implies that the average benefit / average variable cost was $94.50 / 45.74 = 1.82$.

We view this benefit-cost calculation as suggestive rather than conclusive because, like all such calculations, it requires some simplifying assumptions. Most obviously, we use the conservative approach to estimating lifetime benefits described above, we omit average fixed costs because these are very dependent on the scale of the assessment service, and we do not consider general equilibrium effects or the benefits that might be accrued from alternative ways of spending this money. However, we do note that our measure of average variable costs is relatively broad, as it includes semi-fixed costs such as facility and equipment rental, most staff costs, and assessment and software licenses. The only costs that we exclude are those of actually creating the organization, senior management time, and general functions such as accounting.

Finally, we note that there is scope to run similar interventions far more cheaply through job search and matching platforms that incorporate assessments and personalized, automated feedback on assessment results. This approach would reduce or eliminate most of the large components of the average variable cost of in-person assessment: facility rental (27% of average variable cost), equipment rental (22%), and transport (21%).

H Search Effort Appendix

This appendix provides a more detailed description and interpretation of the search effort results and conceptual framework summarized in Section 5. Jobseekers are, on average, overconfident about their skill relative to the reference population: in the tight experiment control group, 63% of beliefs about skill quintiles are above assessment results and only 19% are below; these shares are 46% and 15% in the big experiment (Table C.8). Thus, jobseekers receive on average negative news about their skill levels. If jobseekers react to this information by changing their search effort, this could also affect labor market outcomes.

Treatment effects on beliefs about skill levels: We first document that, on average, jobseekers negatively update their beliefs about their skill levels. In the tight experiment, treatment reduces jobseekers' average beliefs about their level of skills by an average of 0.08 quintiles or 0.15 standard deviations over the two skills (Table H.1, column 1, $p = 0.025$). In the big experiment, treatment reduces jobseekers' average beliefs about their level of skills by an average of 0.11 terciles or 0.3 standard deviations over the three skills (Table H.1, column 3, $p < 0.001$).

Treatment effects vary by baseline beliefs about skill levels, as we might expect. We estimate:

$$Y_{id} = T_d \cdot \beta_1 + T_d \cdot \text{confidence}_i \cdot \beta_2 + \text{confidence}_i \cdot \beta_3 + \mathbf{X}_{id} \cdot \Gamma + \varepsilon_i \quad (6)$$

Where confidence_i is the believed skill level minus the assessed skill level divided by the control group standard deviation.³⁶ We find that $\hat{\beta}_2 < 0$, so treatment effects on belief levels are more negative for jobseekers with higher levels of baseline confidence (Table H.1, row 2). One standard deviation higher baseline confidence is associated with a treatment effect that is 0.26 quintile more negative in the tight experiment (column 2, $p < 0.001$) and a treatment effect that is 0.125 terciles more negative in the big experiment (column 4, $p < 0.001$).

Conceptual framework: These updated beliefs about skill levels might influence search effort. To model this, we replace the assumption of fixed total search effort $\bar{E}$ from Section 2.1 with the assumption that jobseekers choose the levels of search for both communication-heavy and numeracy-heavy jobs, E_C and E_N . This gives a utility function with three arguments: the expected outcome of search for communication-heavy jobs, $V_C(S_C, S_N, E_C)$,

³⁶In the tight experiment, this uses quintiles and averages over two skills; in the big experiment this uses terciles and averages over three skills. We use the term “confidence” to refer to jobseekers' beliefs about their level of skill relative to other jobseekers, not the precision of their beliefs, which we do not measure at baseline. Some researchers refer to this type of overconfidence as “overplacement” (Moore & Healy, 2008).

the expected outcome of search for numeracy-heavy jobs, $V_N(S_C, S_N, E_N)$, and a constraint function capturing the alternative use of time or money allocated to search effort, $A(E_C, E_N)$. For simplicity, we discuss the case of a monetary constraint: $A = Y - P \cdot E_C - P \cdot E_N$, where Y is the jobseeker's unearned income, and P is the price of job search relative to a numeraire consumption good. But we could instead use a time constraint $A = T - E_C - E_N$, where T is the jobseeker's time endowment; a constraint function incorporating time and money; or an intertemporal budget constraint. The jobseeker's problem becomes

$$\max_{E_C, E_N} U(V_C(S_C, S_N, E_C), V_N(S_C, S_N, E_N), Y - P \cdot E_C - P \cdot E_N). \quad (7)$$

As in Section 2.1, we assume that utility is an increasing concave function of all three arguments, that expected search outcomes are increasing concave functions of skill and search effort, and that search effort and skill are more complementary within than across dimensions.

In this framework, increasing the believed level of either skill has an ambiguous effect on total search effort. To see this, note that a fall in the believed level of communication skill S_C lowers the expected marginal productivity of search for communication-heavy jobs, $\frac{\partial V_C}{\partial E_C}$. This has two effects. First, a substitution effect, which causes the jobseeker to substitute away from search for communication-heavy jobs and toward both search for numeracy-heavy jobs and alternative activities. Second, an income effect: it lowers the expected outcome from any given level of search effort, so the jobseeker has to increase search for either communication- or numeracy-heavy jobs to maintain the same expected income. The net effect is a increase in search for numeracy-heavy jobs and an ambiguous effect on search for communication-heavy jobs, and hence an ambiguous effect on total search effort.³⁷

Treatment effects on search effort: We find little evidence that treatment affects search effort in either experiment. In the tight experiment, treatment effects are negative on five of our six search effort measures but all effects are small – less than 11% of the control group mean – and none is statistically significant.³⁸ Treatment lowers an index of these

³⁷The framework has a similar structure to the standard static labor supply model. In that model, a lower wage decreases work effort because the return to work is lower (substitution effect) but increases work effort to afford the same consumption level (income effect). Abebe et al. (2022) also show that raising expected job search outcomes has an ambiguous effect on search effort using a frictional matching model.

³⁸Our six search effort measures are planned applications in the month after the workshop, asked in surveys during the job search workshop; time spent drafting a cover letter during the workshop; click rate on three text messages with links to job adverts sent after the workshop; and three measures of job search on the SAYouth.mobi platform in the month after the workshop: days active, jobs viewed, and applications submitted. The planned applications, text messages, and job applications are described in Section 3.4. The cover letter is a task-based measure of real search effort: the time jobseekers choose to spend drafting a

search effort measures by 0.08 standard deviations (Table H.2, column 1, $p = 0.47$) and a prespecified index of search effort on the SAYouth.mobi platform by 0.1 standard deviations (Table H.3, column 1, $p = 0.29$). In the big experiment, treatment has a tiny effect of 0.003 standard deviations on an index of search effort measures (Table H.4, column 1, $p = 0.92$). Treatment effects on the three components of this index – applications submitted, hours and money spent searching – are positive but tiny ($< 3\%$ of the control group mean) and none is statistically significant.

Treatment effects on search effort also do not vary substantially by jobseekers' baseline confidence levels in either experiment. We estimate equation (6) and show heterogeneous treatment effects by baseline confidence about skills in even-numbered columns in Tables H.2 and H.3 for the tight experiment and Table H.4 for the big experiment.

In the tight experiment, none of the interaction terms are statistically significant, the effect sizes are mostly small, and the signs of the interaction terms vary across search effort measures. The interaction effect on the main search effort index for the tight experiment is a tiny 0.02 (Table H.2, column 2, $p = 0.87$). This implies that a jobseeker with a one standard deviation higher confidence level at baseline has just a 0.02 standard deviation higher treatment effect on search effort. The platform-based measure has a somewhat larger interaction effect of 0.14 standard deviations but it is still not statistically significant (Table H.3, column 2, $p = 0.25$). The key search direction results are also robust to including the interaction with baseline confidence and baseline comparative advantage beliefs in the same regression (Table E.7).

Similarly, treatment effects on search effort in the big experiment do not vary substantially by jobseekers' baseline confidence levels. For the search effort index, the interaction term is a tiny 0.03 standard deviations (Table H.4, column 2, $p = 0.41$). The interaction effects on the index components are positive but small and not statistically significant.

In addition to adjusting search effort, jobseekers who receive negative news about their skills could also redirect their search effort to jobs that they believe are less desirable and hence less competitive. However, we find no evidence that jobseekers choose jobs with different salary levels in the job choice task in the tight experiment (Table H.5). Treatment effects on the assessed salaries of chosen jobs are small (less than 1% of the control mean), far from statistically significant, and do not substantially vary by baseline confidence about skills. This holds when using expert assessments of salaries (columns 1 and 2) and average beliefs of control group jobseekers (columns 3 and 4).

Treatment effects on labor market outcomes by baseline confidence: Finally, we

cover letter for a real job application at the end of the workshop, on a tablet we provided, rather than collecting their incentive and leaving early.

show that treatment effects on labor market outcomes in the big experiment do not vary substantially by baseline confidence about skills (Table H.6). The interaction effects are 0.001 on the work quantity index (column 1, $p = 0.97$) and -0.014 on the work quality index (column 6, $p = 0.70$). The effects on all index components are small and not statistically significant.

Conclusion: This analysis suggests that search effort is unaffected by treatment and hence is unlikely to explain the treatment effects on labor market outcomes. This might arise because the negative treatment effect on believed skill level produces offsetting substitution and income effects on search effort. This does not, of course, imply that search effort plays no role in determining labor market outcomes in this or other settings. See Abebe et al. (2022), Bandiera et al. (2023), and Banerjee & Sequeira (2023), and Mueller & Spinnewijn (2023) for mixed results about the sign of the relationship between search effort and beliefs about labor market prospects.

Table H.1: Heterogeneous Treatment Effects on Skill Beliefs by Baseline Confidence - Both Experiments

	Tight experiment		Big experiment	
	Average skill quintile beliefs (0-4)		Average skill tercile beliefs (0-2)	
	(1)	(2)	(3)	(4)
Treatment	-0.082** (0.035)	0.157* (0.083)	-0.108*** (0.010)	-0.007 (0.011)
Treatment $\times$ Average confidence (bl)		-0.263*** (0.060)		-0.125*** (0.012)
Average confidence (bl)		0.512*** (0.071)		-0.064*** (0.016)
Control mean	2.693	2.693	1.446	1.446
Observations	278	278	4195	4131

Notes: Table H.1 shows that treatment effects on jobseekers' beliefs about their level of skill relative to other jobseekers depend on jobseekers' baseline confidence about their relative skills. Baseline ("bl") confidence refers to jobseekers' baseline belief about their level of skill relative to other jobseekers, minus their assessment result relative to other jobseekers, with a detailed definition in footnote 36. Columns indicate different outcome variables: average belief about communication and numeracy skill quintiles from the tight experiment (cols. 1–2) and average belief about communication, numeracy, and concept formation terciles from the big experiment (cols. 3–4). Even-numbered columns show treatment effects by baseline ("bl") confidence. Control variables are listed in footnotes 15 and 20. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table H.2: Heterogeneous Treatment Effects on Search Effort - Tight Experiment

	Index		Planned apps (w)		Drafting time(w)		SMS click rate		# apps (w)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment	-0.077 (0.103)	-0.101 (0.144)	-3.854 (2.555)	-3.574 (4.124)	-0.530 (0.591)	-0.377 (0.925)	0.003 (0.032)	-0.055 (0.053)	-1.446 (1.544)	-2.068 (2.375)
Treatment $\times$ Avg. confidence (bl)		0.018 (0.111)		-0.429 (2.936)		-0.205 (0.771)		0.062 (0.045)		0.603 (2.058)
Avg. confidence (bl)		0.197 (0.124)		4.032 (2.675)		1.280 (0.900)		-0.059 (0.046)		1.251 (1.907)
Control mean	0.000	0.000	37.878	37.878	8.828	8.828	0.635	0.635	15.194	15.194
Observations	278	278	278	278	267	267	278	278	278	278

Notes: Table H.2 shows that treatment effects on search effort in the tight experiment do not vary significantly by baseline confidence levels. Baseline (“bl”) confidence refers to jobseekers’ baseline belief about their level of skill relative to other jobseekers, minus their assessment result relative to other jobseekers, with a detailed definition in footnote 36. Columns indicate different outcome variables: an inverse covariance-weighted average of the search effort measures used in all other columns (cols. 1–2), the number of planned applications in the 30 days after the workshop (cols. 3–4), the time spent drafting a cover letter during the workshop in minutes (cols. 5–6), the click rate for three SMS messages with links to job adverts we sent to jobseekers after the workshop (cols. 7–8), and the number of applications sent on the job search platform in the 30 days after the workshop (cols. 9–10). Winsorized variables (w) are winsorized at the 99th percentile. Control variables are listed in footnote 15. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table H.3: Heterogeneous Treatment Effects on Platform Search Effort - Tight Experiment

	Platform search effort index		Days active		Adverts clicked (w)		Observed apps (w)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	-0.101 (0.093)	-0.233 (0.158)	-0.723 (0.477)	-1.451** (0.598)	-0.183 (1.530)	-2.692 (3.157)	-1.446 (1.544)	-2.068 (2.375)
Treatment $\times$ Average confidence (bl)		0.138 (0.122)		0.759 (0.449)		2.648 (2.602)		0.603 (2.058)
Average confidence (bl)		-0.061 (0.112)		-0.278 (0.530)		-2.123 (2.304)		1.251 (1.907)
Control mean	-0.000	-0.000	6.014	6.014	9.050	9.050	15.194	15.194
Observations	278	278	278	278	278	278	278	278

Notes: Table H.3 shows that treatment effects on prespecified search effort measures on the job search platform SAYouth.mobi do not vary significantly by baseline confidence levels in the tight experiment. Baseline (“bl”) confidence refers to jobseekers’ baseline belief about their level of skill relative to other jobseekers, minus their assessment result relative to other jobseekers, with a detailed definition in footnote 36. Columns indicate different outcomes: an inverse covariance-weighted average of the search effort measures used in all other columns (cols. 1–2), the number of days jobseekers were active on the platform in the 30 days following the workshop (cols. 3–4), the number of job adverts jobseekers clicked on on the platform in the 30 days following the workshop (cols. 5–6) and the number of applications sent on the job search platform in the 30 days following the workshop (cols. 7–8). Winsorized variables (w) are winsorized at the 99th percentile. Control variables are listed in footnote 15. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table H.4: Heterogeneous Treatment Effects on Search Effort by Baseline Confidence - Big Experiment

	Index		Applications (w)		Hours spent searching (w)		Search expenditure (w)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	0.003 (0.032)	-0.017 (0.038)	0.065 (0.487)	0.021 (0.571)	-0.319 (0.319)	-0.600 (0.501)	0.227 (0.819)	-0.101 (0.857)
Treatment $\times$ Average confidence (bl)		0.030 (0.036)		0.137 (0.464)		0.355 (0.491)		0.504 (0.677)
Average confidence (bl)		-0.034 (0.049)		-0.069 (0.677)		-0.603 (0.657)		-0.428 (0.807)
Control mean	-0.000	-0.000	11.716	11.716	11.083	11.083	20.878	20.878
Observations	4205	4131	4184	4111	4198	4124	4196	4122

Notes: Table H.4 shows that treatment effects on search effort in the big experiment do not vary significantly by baseline confidence levels. Baseline (“bl”) confidence refers to jobseekers’ baseline belief about their level of skill relative to other jobseekers, minus their assessment result relative to other jobseekers, with a detailed definition in footnote 36. Columns show different outcomes: an inverse covariance-weighted average of the search effort measures used in all other columns (cols. 1–2), the number of applications in the 30 days before the endline survey (cols. 3–4), the number of hours spent searching for jobs in the same 30 days (cols. 5–6), and job search expenditure in the same 30 days (cols 7–8). Winsorized variables (w) are winsorized at the 99th percentile. Control variables are described in footnote 20. All monetary values are measured in 2021 USD in purchasing power parity terms. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table H.5: Heterogeneous Treatment Effects on Vertical Search Direction - Tight Experiment

	Average monthly salary of job chosen in job choice task			
	Experts’ assessment		Control group’s assessment	
	(1)	(2)	(3)	(4)
Treatment	2.583 (2.228)	3.655 (3.251)	-0.931 (3.791)	-0.976 (5.968)
Treatment $\times$ Average confidence (bl)		-1.008 (2.242)		-0.068 (4.804)
Average confidence (bl)		-2.495 (2.248)		3.226 (5.789)
Control mean	661.434	661.434	845.472	845.472
Observations	278	278	278	278

Notes: Table H.5 shows that treatment does not lead jobseekers’ to apply to jobs with lower (expected) wages, even for jobseekers’ who receive more negative information about their skills from treatment. Baseline (“bl”) confidence refers to jobseekers’ baseline belief about their level of skill relative to other jobseekers, minus their assessment result relative to other jobseekers, with a detailed definition in footnote 36. The outcome in all columns is the average monthly salary (in 2021 USD in purchasing power parity terms) of the jobs chosen in the job choice task. The salaries are assessed by the hiring experts (cols. 1–2) or by jobseekers in the control group (cols. 3–4). Control variables are listed in footnote 15. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table H.6: Heterogeneous Treatment Effects on Labor Market Outcomes by Confidence - Big Experiment

	Work quantity					Work quality			
	Job offers (w) (1)	Worked month 1 (2)	Worked month 2 (3)	Worked last 7 days (4)	Index (5)	Earnings (w) (6)	Hourly wage (w) (7)	Written contract (8)	Index (9)
Treatment	0.025 (0.022)	0.009 (0.019)	0.015 (0.019)	0.010 (0.017)	0.048 (0.036)	6.461* (3.456)	0.290* (0.172)	0.023* (0.013)	0.097** (0.045)
Treatment $\times$ Average confidence	-0.009 (0.018)	0.022 (0.016)	-0.011 (0.016)	-0.000 (0.014)	0.001 (0.025)	0.100 (2.755)	0.008 (0.132)	-0.007 (0.011)	-0.014 (0.036)
Average confidence	0.019 (0.023)	-0.006 (0.018)	-0.005 (0.020)	-0.019 (0.017)	-0.004 (0.038)	-4.044 (3.640)	-0.016 (0.177)	-0.024 (0.015)	-0.051 (0.046)
Control mean	0.182	0.465	0.437	0.309	-0.000	25.424	1.267	0.120	0.000
Observations	4071	4127	4130	4130	4131	4122	4111	4111	4132

Notes: **Table H.6 shows that treatment effects on labor market outcomes in the big experiment do not vary by baseline confidence levels.** Baseline (“bl”) confidence refers to jobseekers’ baseline belief about their level of skill relative to other jobseekers, minus their assessment result relative to other jobseekers, with a detailed definition in footnote 36. Outcomes are the same as in Table 5: the winsorized number of job offers in the preceding 30 days (col. 1), a dummy indicating any work for pay 1 month after treatment (col. 2), a dummy indicating any work for pay 2 months after treatment (col. 3), a dummy indicating any work for pay in the preceding seven days (col. 4), an inverse covariance-weighted average index of the four employment quantity measures shown in cols. 1–4 (col. 5), winsorized earnings in the preceding seven days (col. 6), winsorized hourly wages in the preceding seven days (col. 7), a dummy indicating a written contract (col. 8), and an inverse covariance-weighted average of the three employment quality measures shown in cols. 6–8 (col. 9). Winsorized variables (w) are winsorized at the 99th percentile. Control variables are listed in footnote 20. All monetary values are measured in 2021 USD purchasing power parity terms. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

I Additional Mechanism Tables

This appendix shows results related to the “additional mechanisms” discussed in Sections 6 and general equilibrium considerations discussed in Section 7.

Table I.1 shows that treatment effects on self-esteem and education investments in the big experiment are very small and not statistically significant. Table I.2 shows treatment effects on the three different willingness-to-pay (WTP) measures in the tight experiment. Treatment lowers jobseekers’ average WTP for a numeracy workbook but does not affect average WTP for a communication workbook. These results suggest that treatment effects on labor market outcomes are unlikely to be explained by effects on educational investment. Table I.2 also shows that treatment does not affect jobseekers’ WTP for information about jobs’ skill demands, suggesting jobseekers seeking this type of information is unlikely to explain effects on labor market outcomes. Further details of the WTP protocol are available at <https://bit.ly/3E34oYH>.

Table I.3 shows that treatment effects on labor market outcomes are driven by jobseekers who did not attach their skill report with any job applications, suggesting that labor market effects are unlikely to be driven by firms learning about assessment results.

Table I.4 shows that the treatment did not induce congestion effects, i.e., it did not change the share of applications directed to numeracy-heavy relative to communication-heavy jobs. This suggests that interventions providing information about skill comparative advantage can help different types of jobseekers apply to different types of jobs, rather than concentrating applications toward one type of job.

Table I.1: Treatment Effects on Additional Mechanisms - Big Experiment

	Self-esteem				Education investment		
	SMS (z) (1)	SMS above med. (2)	Endline (z) (3)	Endline above med. (4)	Any (5)	Apprenticeship (6)	Formal (7)
Treatment	0.003 (0.008)	0.014 (0.014)	0.007 (0.021)	0.014 (0.015)	0.011 (0.011)	0.006 (0.005)	0.009 (0.011)
Control mean	-0.000	0.483	0.000	0.471	0.224	0.036	0.185
Observations	3334	3334	4206	4206	4205	4205	4205

Notes: Table I.1 shows that treatment does not affect average self-esteem or education investment in the big experiment. Cols. 1–2 show effects on self-esteem in the SMS survey 2–3 days after the workshops. Cols. 3–4 show effects on self-esteem in the endline survey 3.5 months after the workshops. The SMS and phone surveys use respectively one and five items from the Rosenberg (1965) scale, both answered on five-point Likert scales. Cols. 1 and 3 use standardized measures and columns 2 and 4 use dummies for above-median values. Cols. 5, 6, and 7 show effects on dummies for enrolling in respectively any education, an apprenticeship, and a formal degree/diploma. Control variables are listed in footnote 20. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table I.2: Treatment Effects on Willingness-to-pay - Tight Experiment

	Info on skill requirements						Numeracy materials						Communication materials					
	Pooled		Num. CA		Comm. CA		Pooled		Num. CA		Comm. CA		Pooled		Num. CA		Comm. CA	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
Treatment	-0.719 (5.766)	-6.851 (11.255)	-12.823 (13.609)	-14.035 (15.888)	1.708 (8.396)	-8.179 (17.726)	-13.410*** (3.539)	-15.309** (6.775)	-14.527* (8.049)	-20.583** (8.944)	-10.046** (3.898)	-3.217 (11.201)	-0.822 (3.790)	-0.247 (6.747)	5.252 (9.785)	-2.571 (12.162)	-3.821 (5.882)	3.862 (11.719)
Treatment × Aligned CA belief (bl)		12.287 (18.773)		8.804 (38.367)		14.077 (22.913)		5.019 (11.915)		36.732** (15.806)		-9.825 (16.400)		-0.847 (10.474)		45.821** (21.394)		-11.987 (14.833)
Aligned CA belief (bl)		0.174 (16.202)		-12.737 (28.007)		7.989 (16.687)		-14.744 (8.948)		-38.444*** (11.399)		-4.144 (11.074)		-3.753 (9.163)		-43.991*** (11.637)		7.831 (10.853)
Treatment effect: Aligned CA belief (bl)		5.436 (10.847)		-5.231 (33.252)		5.898 (11.314)		-10.290 (7.268)		16.149 (14.781)		-13.042* (7.247)		-1.094 (6.130)		43.251** (16.512)		-8.125 (7.682)
Control mean	78.867	78.867	79.337	79.337	78.611	78.611	54.964	54.964	55.408	55.408	54.722	54.722	48.327	48.327	49.184	49.184	47.861	47.861
Observations	277	277	105	105	172	172	277	277	105	105	172	172	277	277	105	105	172	172

Notes: Table I.2 shows treatment effects on willingness-to-pay (WTP) for different products relevant to job search in the tight experiment. “CA” stands for comparative advantage in skills and “bl” stands for baseline. Cols 1–6 show effects on the WTP for a document with the expert-assessed skill requirements for the 11 job pairs in the job choice task. Cols. 7–12 show effects on WTP for a numeracy training resource. Cols. 13–18 show effects on WTP for a communication training resource. Cols. 1–2, 7–8, and 13–14 show results for the full sample. Cols. 3–4, 9–10, and 15–16 show results for jobseekers with a comparative advantage in numeracy. Cols. 5–6, 11–12, and 17–18 show results for jobseekers with a comparative advantage in communication. All currency values are in 2022 South African Rands, which was the currency used in the actual WTP measurement exercise and is consistent with the supplemental WTP appendix posted at <https://bit.ly/3E34oYH>. Multiplying these values by 0.140 converts them into the same currency units as the rest of the paper, 2021 USD in purchasing power parity terms. Control variables are listed in footnote 15. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table I.3: Heterogeneous Treatment Effects on Labor Market Outcomes by Skill Report Attachment to Applications - Big Experiment

	Work quantity					Work quality			
	Job offers (w) (1)	Worked month 1 (2)	Worked month 2 (3)	Worked last 7 days (4)	Index (5)	Earnings (w) (6)	Hourly wage (w) (7)	Written contract (8)	Index (9)
Treatment	0.040 (0.027)	0.033** (0.016)	0.018 (0.018)	0.014 (0.016)	0.068* (0.039)	8.572*** (3.020)	0.359** (0.143)	0.023** (0.011)	0.110*** (0.039)
Treatment × Attached report w. application	-0.038 (0.038)	-0.021 (0.022)	-0.039 (0.024)	-0.013 (0.024)	-0.066 (0.049)	-5.108 (4.229)	-0.179 (0.211)	-0.019 (0.017)	-0.071 (0.058)
Control mean	0.195	0.465	0.437	0.309	-0.000	25.424	1.267	0.120	0.000
Observations	3933	3988	3991	3991	3992	3983	3971	3971	3993

96

Notes: Table I.3 shows that treatment effects on labor market outcomes are driven by jobseekers that do not use their skill reports in any job applications in the big experiment. These results should be interpreted with caution because the right-hand side of the regression includes the interaction between treatment and a post-treatment outcome, “Attached report w. application.” The interaction term is included in the regression but the variable “Attached report w. application” is omitted from the regressions because no control individual received a report. Outcomes are the same as in Table 5: the winsorized number of job offers in the preceding 30 days (col. 1), a dummy indicating any work for pay 1 month after treatment (col. 2), a dummy indicating any work for pay 2 months after treatment (col. 3), a dummy indicating any work for pay in the preceding seven days (col. 4), an inverse covariance-weighted average index of the four employment quantity measures shown in cols. 1–4 (col. 5), winsorized earnings in the preceding seven days (col. 6), winsorized hourly wages in the preceding seven days (col. 7), a dummy indicating a written contract (col. 8), and an inverse covariance-weighted average of the three employment quality measures shown in cols. 6–8 (col. 9). Winsorized variables (w) are winsorized at the 99th percentile. Control variables are listed in footnote 20. All monetary values are measured in 2021 USD purchasing power parity terms. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table I.4: Treatment Effects on Concentration of Job Search - Tight Experiment

	Degree of job search concentration		
	Individual level measure		Job-pair level measure
	(1)	(2)	(3)
Treatment	-0.007 (0.017)	0.009 (0.019)	-0.029** (0.013)
Treatment $\times$ Aligned CA belief (bl)		-0.032 (0.028)	
Aligned CA belief (bl)		0.050** (0.024)	
Treatment effect: Aligned CA belief (bl)		-0.023 (0.024)	
Control mean	0.181	0.181	0.077
Observations	278	278	22

Notes: **Table I.4 shows that the treatment weakly decreases the concentration of jobseekers' applications in the job choice task.** Cols. 1 and 2 show effects on the concentration of job choices in the job choice task using one observation per jobseeker. This concentration measure is constructed as the absolute deviation of the fraction of chosen numeracy jobs from 0.5 averaged across job pairs at the jobseeker level. Col. 3 shows effects on the concentration of job choices in the job choice task using one observation per job pair. This concentration measure is constructed in two steps. First, we calculate the fraction of jobseekers choosing the numeracy job in each job-pair $\times$ treatment group combination. Second, we calculate the absolute deviation of this measure from 0.5. For both concentration measures, higher numbers indicate a higher degree of concentration of job choices. We do not estimate heterogeneous treatment effects by baseline aligned comparative advantage belief for the job-pair level analysis in col. 3 because the outcome of interest is averaged across jobseekers. Controls used in cols. 1 and 2 are listed in footnote 15. Standard errors are clustered at the treatment-day level in cols. 1 and 2 and are heteroskedasticity-robust in col. 3. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

J Gender

This appendix reports descriptive statistics and treatment effects separately for women and men. Tables J.1 and J.2 display gender differences in baseline beliefs about skills in respectively the tight and big experiments. Gender differences in beliefs are small and are generally not statistically significant after adjusting for gender differences in demographics, education, and assessment results. Tables J.3 and J.4 display treatment effects on skill beliefs by gender in respectively the tight and big experiments. The treatment effects on skill beliefs do not differ by gender. The lack of gender differences in treatment effects on beliefs about skills makes gender differences in treatment effects on downstream outcomes (e.g. search direction, job quality) unlikely, so we do not report gender-disaggregated treatment effects on these outcomes.

Table J.1: Gender Differences in Beliefs about Skills - Tight Experiment

	Female	Male	Δ	$p(\Delta = 0)$	$\Delta(\text{adjusted})$	$p(\Delta(\text{adjusted}) = 0)$	N
Aligned CA belief	0.50	0.46	-0.04	0.53	-0.01	0.83	278
Fraction aligned beliefs	0.23	0.21	-0.02	0.46	-0.05	0.15	278
Fraction overconfident beliefs	0.62	0.59	-0.03	0.58	0.04	0.11	278
Fraction underconfident beliefs	0.15	0.20	0.05	0.12	0.01	0.79	278

Notes: Table J.1 shows that gender differences in baseline beliefs about skills in the tight experiment are small. Adjusted differences control for prespecified covariates described in footnote 15, except the baseline value of the outcome variable. “CA” stands for comparative advantage in skills. P-values are estimated from regressions that use heteroskedasticity-robust standard errors.

Table J.2: Gender Differences in Beliefs about Skills - Big Experiment

	Female	Male	Δ	$p(\Delta = 0)$	$\Delta(\text{adjusted})$	$p(\Delta(\text{adjusted}) = 0)$	N
Aligned CA belief	0.19	0.23	0.03	0.01	0.03	0.02	4312
Fraction aligned beliefs	0.35	0.42	0.07	0.00	0.01	0.19	4378
Fraction overconfident beliefs	0.53	0.45	-0.08	0.00	0.00	0.73	4378
Fraction underconfident beliefs	0.11	0.13	0.01	0.06	-0.01	0.03	4378

Notes: Table J.2 shows that gender differences in baseline beliefs about skills in the big experiment are small. Adjusted differences control for prespecified covariates described in footnote 20, except the baseline value of the outcome variable. “CA” stands for comparative advantage in skills. P-values are estimated from regressions that use heteroskedasticity-robust standard errors.

Table J.3: Heterogeneous Treatment Effects on Beliefs about Skills by Gender - Tight Experiment

	Aligned CA belief		Fraction aligned beliefs	
	(1)	(2)	(3)	(4)
Treatment	0.181*	0.178**	0.087	0.063
	(0.091)	(0.079)	(0.079)	(0.051)
Treatment $\times$ Female	-0.074	-0.062	0.059	0.022
	(0.107)	(0.103)	(0.095)	(0.062)
Female	0.082	0.019	-0.038	-0.033
	(0.066)	(0.063)	(0.060)	(0.042)
Treatment effect: Female	0.107	0.116**	0.146***	0.085**
	(0.065)	(0.047)	(0.049)	(0.032)
Control mean	0.475	0.475	0.183	0.183
Observations	278.000	278.000	278.000	278.000
Controls	No	Yes	No	Yes

Notes: Table J.3 shows that treatment effects on skill beliefs do not substantially differ by gender in the tight experiment. “CA” stands for comparative advantage in skills. Columns 1 and 2 show effects on a dummy equal to one if jobseekers’ beliefs about their skill comparative advantage are aligned with their assessment results. Columns 3 and 4 show treatment effects on the fraction of skill beliefs that align with measured skill quintiles. Control variables are listed in footnote 15. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table J.4: Heterogeneous Treatment Effects on Beliefs about Skills by Gender - Big Experiment

	Aligned CA belief		Fraction aligned beliefs	
	(1)	(2)	(3)	(4)
Treatment	0.156***	0.155***	0.150***	0.153***
	(0.021)	(0.020)	(0.019)	(0.015)
Treatment $\times$ Female	-0.032	-0.026	-0.021	-0.016
	(0.027)	(0.026)	(0.020)	(0.017)
Female	-0.029	-0.011	-0.051***	0.003
	(0.018)	(0.016)	(0.015)	(0.012)
Treatment effect: Female	0.124***	0.130***	0.128***	0.136***
	(0.013)	(0.015)	(0.013)	(0.010)
Control mean	0.196	0.196	0.388	0.388
Observations	4191	4118	4205	4195
Controls	No	Yes	No	Yes

Notes: Table J.4 shows that treatment effects on skill beliefs do not substantially differ by gender in the big experiment. “CA” stands for comparative advantage in skills. Columns 1 and 2 show effects on a dummy equal to one if jobseekers’ beliefs about their skill comparative advantage are aligned with their assessment results. Columns 3 and 4 show treatment effects on the fraction of skill beliefs that align with measured skill quintiles. Control variables are listed in footnote 20. Standard errors clustered at the treatment-day level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

K Preregistration Appendix

The paper uses two preregistrations and analysis plans: [AEARCTR-0001631](#) for the big experiment and [AEARCTR-0010000](#) for the tight experiment. The preanalysis plan for the big experiment covers the experimental design, estimating equation for the average treatment effects, covariates, inference methods, and labor market outcomes. It does not include outcomes on comparative advantage beliefs and skill-directed search or the heterogeneity analysis by aligned baseline comparative advantage beliefs.

The tight experiment was set up to test the research question that arose from the exploratory analysis of the big experiment. The research question, experimental design, estimating equations, covariates, inference methods, and outcomes are prespecified. The heterogeneity analyses by aligned baseline comparative advantage beliefs and baseline confidence about levels of skills are prespecified as additional analyses. We depart from the pre-analysis plan for the tight experiment in four ways:

1. We restrict the sample to jobseekers who have a clear comparative advantage because the skill-directed search measures can only be sensibly defined for these jobseekers. When we use this restricted sample, we omit one prespecified control variable, the dummy variable about whether the jobseeker has a clear comparative advantage, because this control variable has no variation in the restricted sample. In Appendix Tables [E.1](#) and [E.2](#) we show the main results of the paper for the full sample as a robustness check. The interpretation of results remains unchanged.
2. We added further skill-directed search outcomes: clicks on jobs sent in SMS messages and job search on the SAYouth.mobi platform. We obtained these measures from our partner organization after we had lodged the pre-analysis plan. We correct for the inclusion of these additional measures by creating a summary index of all skill-directed search outcomes.
3. To align the search direction and search effort measures, we add the SMS click rate and the number of observed application clicks to the main search effort table and, again, construct a summary index. We show results for the prespecified platform search index in Table [H.3](#).
4. Following recent methodological critiques ([Chen & Roth, 2023](#); [Mullahy & Norton, 2022](#)) we use winsorization instead of the prespecified inverse hyperbolic sine transformation for outcomes with potentially long right tails (e.g. earnings, number of applications). Using the inverse hyperbolic sine transformation produces qualitatively similar results, which we show for earnings in Table [E.5](#).