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

Andrea Kiss ⓡ Robert Garlick ⓡ Kate Orkin ⓡ Lukas Hensel

June 8, 2023

Abstract

Worker sorting into tasks and occupations based on their skills plays a potentially important role in aggregate labor productivity. This sorting may be inefficient if jobseekers do not apply to jobs that match their skills. We test this idea by studying how jobseekers' beliefs about their skills affect their job search activities and labor market outcomes. We assess communication and numeracy skills of young South African jobseekers to identify their comparative advantage across these skills. Self-perceived and measured comparative advantage persistently differ for many of these jobseekers. In a field experiment, giving jobseekers their skill assessment results reduces differences between perceived and measured skills. This leads jobseekers to redirect search toward jobs whose skill demand matches their comparative advantage in multiple search measures: an incentivized job choice task where workers choose between jobs with different skill requirements, detailed usage data from an online job search platform, and self-reported search plans. In a larger field experiment with a longer follow-up period, treatment has similar effects on self-reported measures of beliefs and search direction and substantially raises earnings and job quality, although not employment rates. These patterns are consistent with models of endogenously directed job search, where jobseekers' beliefs about their skills influence where they direct job applications and hence the wages they are offered.

*Robert Garlick, Duke University, robert.garlick@duke.edu; Lukas Hensel, Guanghua School of Management, Peking University, lukas.hensel@gsm.pku.edu.cn; Andrea Kiss, Carnegie Mellon University, akiss@andrew.cmu.edu; Kate Orkin, Blavatnik School of Government, University of Oxford, kate.orkin@bsg.ox.ac.uk. Order of authors is randomized and all authors contributed equally. For helpful comments, we thank Johannes Abeler, Peter Arcidiacono, Michele Belot, Clement Imbert, David Huffman, Jeremy Magruder, Chris Roth, Duncan Thomas, and Basit Zafar, as well as seminar and conference participants at AFE, Carnegie Mellon, Columbia, Duke, Essex, FIU, IFLAME, Oxford, Peking, PHBS, Renmin, Pittsburgh, and Royal Holloway. We thank Alice Cahill, Raquel Caldeira, Aliya Chikte, Sabhya Gupta, Jenn Kades, Brynde Kreft, and Wim Louw for excellent research assistance. This study has been approved by the University of Cape Town Commerce Faculty Ethics in Research Committee, the Duke Institutional Review Board for the Protection of Human Subjects in Non-Medical Research (# D0368) and the Oxford Department of Economics Research Ethics Committee (# ECONCIA15-055). The experiments used in this study are preregistered on the AEA's trial registry at https://doi.org/10.1257/rct.1631-8.0 and https://doi.org/10.1257/rct.10000-1.0. This paper is 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, and W. E. Upjohn Institute for Employment Research.

1 Introduction

Worker sorting into tasks and occupations has long been recognized as an important feature of the labor market (Roy, 1951). Efficiently matching workers with the tasks where their skills are most productive offers the prospect of substantial gains in output (Lise and Postel-Vinay, 2020). The efficiency of sorting depends crucially on the job search process, where jobseekers choose how to direct their search effort to different types of jobs. However, jobseekers may have imperfect information about their productivity in different types of jobs, causing them to direct search effort to jobs that poorly match their skills.

This paper studies how jobseekers' beliefs about their own skills influence their job search behavior and outcomes. We emphasize how beliefs about relative skill ranking in different dimensions, which we call **comparative advantage beliefs**, influence search across jobs with different skill demands, which we call **skill-directed job search**. We run two experiments with complementary designs. In both, we assess jobseekers' skills in multiple domains, measure their beliefs about these skills, experimentally shift these beliefs by providing assessment results, and measure their post-treatment search activities and outcomes. We find that giving jobseekers assessment results moves their comparative advantage beliefs toward their measured comparative advantage, redirects search to jobs that match their comparative advantage in skills, and leads to substantial increases in their earnings and wages but not their employment rate. These results suggest jobseekers face information frictions that limit their ability to match to jobs where they will be most productive. Existing studies cannot answer this question because they lack the necessary combination of data: jobseeker-level data on multidimensional skills, beliefs about these skills, and job search activities and outcomes; application-level data about skill demands; and exogenous variation in beliefs about skills.[1]

We begin by describing skills, beliefs, and job search of young jobseekers from disadvantaged backgrounds in Johannesburg, South Africa. We assess their skills in communication and numeracy – two general-purpose, job-relevant skills that we verify firms in this context value in prospective workers. We then measure jobseekers' beliefs about their own skills and document substantial differences between measured and perceived skills, both expressed as ordinal rankings. In particular, measured and perceived comparative advantage differ for almost half the jobseekers.[2] These differences persist through time, even for those who find work or search intensively, perhaps because jobseekers report extremely limited feedback on their skills from job applications. Jobseekers' skill beliefs also predict their job search choices and jobseekers believe their search outcomes would improve if their skills were higher.

[1] Survey data typically record aggregated search data such as total applications submitted, not application-level data including skill demand. Job search platform data seldom measure skills, beliefs, or labor market outcomes. Administrative data from unemployment insurance systems do not measure beliefs and seldom measure skills or search.

[2] We define each jobseeker's comparative advantage as the skill dimension in which she ranks highest, following some of the few existing job search studies using multidimensional skill measures (e.g. Guvenen et al. 2020). This definition has both advantages and disadvantages relative to some other approaches in labor economics, as we discuss later in the paper. In particular, it avoids the need for any information about labor demand and avoids the challenging task of estimating the relative prices of different skills in the labor market.

We propose a simple model to capture these patterns and show how these skill beliefs can influence job search. In the model, jobseekers choose the levels of search effort directed to jobs with different skill requirements. Inaccurate skill beliefs distort the perceived return to searching for each type of job and hence distort search decisions. If jobseekers apply to jobs that poorly match their skills, this can reduce the number of job offers received or wages offered.

We then run two field experiments to identify the relationship between comparative advantage beliefs, skill-directed job search, and search outcomes. Our first experiment (N = 278) focuses on skill-directed job search and is embedded in day-long job search workshops run with a job search assistance agency that include taking the skill assessments. We randomly assign half the workshops to treatment: giving participants their skill assessment results. Treated participants' skill beliefs are 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 different measures: applications in a novel, incentivized task in which participants choose between applying to jobs with different skill requirements; planned applications in an end-of-workshop survey; and actual applications on a job search platform in the month after the workshop. These results show the link between comparative advantage beliefs and skill-aligned job search using unusually rich data but do not show if this affects search outcomes.

Our second experiment (N = 4,389) focuses on labor market outcomes. This experiment is also embedded in day-long job search workshops that include skill assessments. We again randomly assign half of the workshops to give participants their skill assessment results. Survey data collected on average 3.5 months after the workshops show that treatment improves labor market outcomes: treated jobseekers' weekly earnings are 6.5 USD higher and hourly wages are 0.3 USD higher than control group jobseekers (23% and 26% of control group mean, respectively) and they are more likely to have formal job contracts, although we find at most weak evidence for a higher employment rate. The average earnings gain between treatment and the follow-up survey is roughly 1.8 times the average variable cost of running the assessment system, suggesting a potential role for these types of interventions in active labor market policies. Treatment effects on jobseekers' comparative advantage beliefs and skill-directed job search are qualitatively similar to the first experiment, although measurement differences make quantitative comparisons difficult.

The two field experiments and the descriptive data show a consistent picture of this labor market: shifting jobseekers' perceived comparative advantage toward their measured comparative advantage redirects their search toward jobs aligned with their measured comparative advantage and leads to better-quality employment. They provide two complementary types of evidence. The first experiment provides unusually rich survey, task, and platform data on beliefs and search, including application-level data. The second experiment provides survey data on beliefs, search, and labor market outcomes for a larger sample over a longer time period. Given the two experiments' different strengths, we call them the "tight" and "big" experiments respectively throughout

the paper.[3] The similar belief and search results across two experiments show some generalizability for the economic process we study because the experiments are conducted six years apart with similar but not identical implementation and similar but not identical samples.

We find limited evidence for several other theory-informed ways that jobseekers might respond to new information about their skills. On average, jobseekers in our sample believe at baseline that their skills are higher than their measured skills, so treated jobseekers receive on average negative news about their skill levels. The effect of lower perceived skill level on search effort and hence labor market outcomes is theoretically ambiguous in our model. In practice, we find negligible effects on search effort for jobseekers on average and for those who receive positive or negative news, suggesting that changes in overall search effort are unlikely to explain the observed effects on labor market outcomes. Furthermore, treatment does not affect jobseekers' generalized self-esteem, just their beliefs about their skills; does not affect their educational investment, at least over the time horizon we observe; and does not provide information directly to firms about jobseekers' skills.

More generally, what labor market conditions might lead to inaccurate comparative advantage beliefs and hence to poorly skill-directed job search? Jobseekers might enter the labor market with inaccurate comparative advantage beliefs in the many countries where education systems give students inaccurate feedback on their performance (Pritchett, 2013). High unemployment might limit jobseekers' opportunities to learn about their skills through work experience, a standard learning mechanism in many models (e.g. Guvenen et al. 2020). And jobseekers might struggle to evaluate their fit with different jobs when navigating new or rapidly changing labor markets due to industrial displacement, migration, or structural transformation. This form of limited information might be particularly costly when screening many mismatched job applicants is costly to firms and when job search is expensive for jobseekers (Abebe et al., 2021a; Banerjee and Sequeira, 2020), particularly if well-matched jobseekers face systematically higher application costs (Abebe et al., 2021b).

This paper makes two primary contributions. First, we provide the first direct evidence that jobseekers' beliefs about their comparative advantage in skills influence how they direct search to different job types and hence influence their search outcomes. Sorting on skill comparative advantage is central to both classic and contemporary labor market models but is less studied empirically (Roy, 1951; Sanders and Taber, 2012). There is a growing literature studying the relationship between jobseekers' beliefs, search activities, and search outcomes, reviewed by Mueller and Spinnewijn (2022). This includes work on the co-evolution of beliefs and search in panel data (Conlon et al., 2018; Mueller et al., 2021; Spinnewijn, 2015); experiments that provide information about labor market conditions (Altmann et al., 2018; Jones and Santos, 2022); and work on search subsidies, matching services, or mentoring programs that influence jobseekers' beliefs

[3]We recognize this is non-standard terminology and hope readers will take this in the lighthearted spirit we intend. Standard terminology is not well-suited to distinguish the two experiments. Both are field rather than lab or lab-in-the-field experiments because they feature real jobseekers engaging in real job search activities. Both have elements of what Harrison and List (2004) call framed field experiments and natural field experiments.

about their labor market prospects (Abebe et al., 2022; Alfonsi et al., 2022; Bandiera et al., 2021; Banerjee and Sequeira, 2020; Kelley et al., 2020).[4] This work focuses on jobseekers' beliefs about the level of their labor market prospects, captured by job offer arrival rates or wage offer distributions. But this does not study beliefs about comparative advantage, perhaps because very few datasets combine information on jobseekers' skills and vacancies' skill demands. We complement this literature by collecting unique data to study beliefs about comparative advantage and how search is directed toward jobs with different skill demands. Comparative advantage beliefs are likely to influence the beliefs about labor market prospects that are more widely studied. Methodologically, our work is closest to Altmann et al. (2022) and Belot et al. (2022a, 2018), who show that encouraging jobseekers to apply to different occupations can improve some labor market outcomes, consistent with jobseekers having imperfect information about their job offer probabilities in some occupations. Like these papers, we study how jobseekers direct search to different job types using jobseeker×application-level data, rather than relying only on more common jobseeker-level measures such as the total number of job applications they submit. Unlike these papers, we study the role of comparative advantage and measure jobseekers' skills, beliefs, and a broader range of search behavior and outcomes.

Second, our work complements a growing literature on the role of job search in worker-firm matching with multidimensional skills (e.g. Lise and Postel-Vinay 2020). Recent work uses dynamic models to estimate large lifetime earnings losses when jobseekers enter the labor market with imperfect information about their skills and hence their productivity in different jobs (e.g. Guvenen et al. 2020). These papers use rich longitudinal data on employment, guided by dynamic models, to show long-term consequences of the belief-induced search frictions that we document. But they do not observe beliefs or search. Our measures of skills, beliefs about skills, search, and short-term labor market outcomes, combined with experimental variation in beliefs, allow us to directly assess the assumptions of these models. Within this literature, our findings are most consistent with the modeling framework of Baley et al. (2022), who study the long-term consequences of jobseekers directing search based on inaccurate beliefs about their multidimensional skills.

A parallel literature focuses on firm-side limited information about job applicants' skills.[5] We differ by focusing on jobseekers' assessments of their own skills. This also matters for firms, because the self-assessments determine how jobseekers direct search, which in turn determines the pool of applicants to each job. Carranza et al. (2022) study the relative roles of firm- and jobseeker-side learning about jobseekers' skills using some data from our "big" field experiment. This paper focuses more deeply on jobseeker-side learning and differs from Carranza et al. (2022) in several

[4]This relates to research into jobseekers' beliefs about attributes of specific jobs (e.g. Bazzi et al. 2021; Heath et al. 2022; Sockin and Sojourner 2020; Subramanian 2022).

[5]For example, providing firms with additional information about prospective workers' skills can increase employment and earnings (Abebe et al., 2021a; Abel et al., 2020; Bassi and Nansamba, 2022; Pallais, 2014), some wage-tenure relationships are consistent with firms learning about current workers' productivity over time (Altonji and Pierret, 2001; Arcidiacono et al., 2010; Kahn and Lange, 2014), and workers have better outcomes in some labor market outcomes when they have formal educational qualifications, conditional on measured skills (Alfonsi et al., 2017; MacLeod et al., 2017).

ways. This paper adds the "tight" experiment, uses more detailed measures of jobseekers' beliefs and search activity from the shared field experiment, introduces the idea of comparative advantage beliefs, and substantially expands the idea of skill-directed job search, both empirically and theoretically. In contrast, Carranza et al. (2022) use additional firm-facing experiments to show how firms respond to new information about jobseekers' skills, which this paper does not study.

Our work also relates to research into the formation and consequences of beliefs about skills in contexts other than job search. Some papers study how skill beliefs influence decisions in the workplace, typically focusing on the role of overconfidence rather than comparative advantage beliefs (e.g. Hoffman and Burks 2020; Huffman et al. 2022; Malmendier and Tate 2015). A growing literature in education shows that potentially noisy self-assessments of skill levels and comparative advantage matter for enrollment and subject choices, even though education might offer scope for more structured feedback on skills than job search (e.g. Arcidiacono et al. 2016; Stinebrickner and Stinebrickner 2014a,b). Finally, a large literature studies the impact of beliefs on decisions in lab settings, including a few papers studying the role of beliefs about skills in lab tasks that mimic job search (e.g. Falk et al. 2006).

Methodologically, we add to a small, recent literature showing the value of combining multiple types of experiments to understand both a job search process and the outcomes of that process (Carranza et al., 2022; Cortés et al., 2021; Field et al., 2023). The job search task in our tight experiment provides a novel way to use jobseekers' choices over job applications to measure their valuation of skill match relative to other vacancy characteristics, building on work using choice experiments to study job choice or labor supply (Mas and Pallais, 2017; Wiswall and Zafar, 2015).

The paper is organized into four sections. Section 2 describes the context and sample, patterns of measured and perceived skills in the sample, and a simple model of job search with potentially inaccurate skill beliefs. In Section 3 we show the relationship between comparative advantage beliefs and skill-directed job search in the tight field experiment. In Section 4 we show the relationship between comparative advantage beliefs, skill-directed job search, and labor market outcomes in the big field experiment. Section 5 discusses the role of additional mechanisms.

2 Economic Environment

This section describes the samples of jobseekers, how we assess their skills, their beliefs about their skill comparative advantage relative to other jobseekers, how these beliefs correlate with their search behavior, and that the skills we focus on are valued by firms and perceived as valuable by jobseekers.

2.1 Sample and Context

We study young, active jobseekers with high school education from low-income backgrounds in Johannesburg, South Africa. We recruited two samples, one of 278 people in 2022 for our tight experiment, and one of 4389 people in 2016-7 for our big experiment, using similar recruitment methods. We recruited both samples from participant registries at the Harambee Youth Employ-

ment Accelerator, a social enterprise funded jointly by the South African government and private firms to improve labor market functioning, including through job search assistance services to young jobseekers from low-income backgrounds.

Sample recruitment: We recruited participants for the tight field experiment from active jobseekers signed up to Harambee's online job search and matching platform SAYouth.mobi. This platform allows firms to list vacancies and jobseekers to apply for jobs and access resources to support job search. Harambee recruits jobseekers to the platform through advertising in traditional and social media. At the time of the study, roughly 3 million jobseekers were registered on the platform. Jobseekers are eligible to use the platform if they self-report being aged 18-34, have legal permission to work in South Africa, have completed secondary school, are from disadvantaged backgrounds, and are actively searching for work. Jobseekers often receive text messages from Harambee with information on employment opportunities, including invitations to take assessments to determine eligibility for types of vacancies.

We contacted a random sample of jobseekers who used the platform at least once in the 30 days before the start of the experiment in July 2022 and who live within commuting distance of our field location in downtown Johannesburg. We invited them to a day-long job search assistance workshop, where they are told they will take skill assessments, answer questions about their job search, and receive job search advice, similar in format to other workshops run by Harambee. We ran our experiment during this workshop. Invitations were sent by text message and have the hashtag #SAYouth, which Harambee commonly uses. Jobseekers who replied to indicate interest received the workshop details and were told they would receive reimbursement for transport and the time spent at the workshop. We ran 34 day-long workshops, each with 8 to 15 jobseekers, for a total of 373 jobseekers. Our final sample is restricted to the 278 jobseekers whose skill assessment results allow us to clearly identify their skill comparative advantage for reasons we explain in Section 2.2.

For the big experiment, we recruited 4,389 participants using a similar approach between September 2016 and April 2017. This predated the launch of SAYouth.mobi but Harambee ran a similar process of recruiting jobseekers and providing job search assistance services. Harambee invited jobseekers to attend day-long workshops in downtown Johannesburg that included some job search advice and standardized skill assessments that could be used to match them to employers with suitable vacancies.[6] We administer the big experiment over 54 skill assessment days, with an average of 80 jobseekers per day.

Baseline measurement: In both experiments, participants take skill assessments and complete surveys about their demographics, beliefs about their skills, beliefs about their job search, recent job search activities, current employment, and employment history before treatment. We use these pre-treatment measures to describe the participants in this section. We leave discussion of post-treatment measurement to sections 3 and 4.

[6]Some candidates were invited to further job readiness training and placement services based on their assessment results and residential location, but only 0.2% of candidates in our sample get jobs through this training, so this is not a central feature for interpreting our results.

Descriptive statistics: We focus here on descriptive statistics for the tight experiment sample, shown in Table 1. These jobseekers are actively engaged in the labor market but likely to face limited information about their skills in the labor market for multiple reasons. First, they are young and so have limited time to learn about their skills from search or employment (interdecile age range 21 - 32). Second, only 23% had completed any post-secondary education, limited their scope to learn from specialized training. Third, they have limited, mostly informal work experience: 33% were employed at baseline but only 13% had a formal written contract and only 25% had *ever* held a long-term wage job.[7] Their job search effort was high but met with limited success. 96% were actively searching and the average jobseeker submitted 10 job applications, spent 14 hours searching for work, and spent 22.72 USD on search expenses in the week before the baseline. This search cost is high – equal to 50% of the average weekly earnings of employed people in our sample. The average jobseeker received only 0.17 job offers in the last 30 days, implying that $< 1\%$ of applications yield offers. (We used a longer recall period for offers than applications due to their rarity.) Importantly, the jobseekers we study are not a population-representative sample, but relevant to the research question we study.

The sample for the big experiment, described in Section 4.1 and Table A22, is very similar. The field experiment recruitment occurred before Harambee set up the online job search platform, potentially contributing to slight differences in sample composition. On search behavior and labor market outcomes, the samples are strikingly similar. But the field experiment sample is slightly younger (24 vs 26 years); slightly better-educated, with 39% having attended post-secondary schooling relative to 23% for the tight experiment; but slightly less experienced, with 9% having held a permanent or long-term job relative to 25% in the tight experiment.

Context: Johannesburg's labor market and our sample have three salient features that limit jobseekers' scope to learn about their skills. First, the education system provides noisy signals of jobseekers' skills. Most of our sample have secondary education as their highest level of qualification and, due to Harambee's eligibility criteria, attended school in a low-income neighborhood. In such schools, grades and grade progression are weakly correlated with results in independent skill assessments (Lam et al., 2011; Taylor et al., 2011). There is a nationally standardized secondary school graduation examination, but grades in this exam only weakly predict performance in post-secondary education and firms report limited trust in how school-leaving scores predict workplace performance (Schoer et al., 2010). Schools in these neighborhoods rarely have specialized career counselors and few teachers have received training to offer support in choice of career (Pillay, 2020).

Second, the employment rate is low, especially for formal jobs. At the time of the tight experiment, unemployment in Johannesburg was 33.7% for the working-age population and 40.5% for ages 15-34 (Statistics South Africa, 2022). In our sample, unemployment is even higher at 67%, most likely because our sample omits people outside the labor force. Jobseekers' limited expe-

[7] Throughout the paper, we define an employed person as someone who did any income-generating activity during the survey reference period, following Statistics South Africa's definition. Unemployment rates exclude those in full-time education or not in the labor force.

Table 1: Summary Statistics

	(1) Mean	(2) Median	(3) Min	(4) Max	(5) SD	(6) Obs.
Panel A: Demographics						
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	0.69	1.00	0.00	1.00	0.46	278
University degree / diploma	0.09	0.00	0.00	1.00	0.28	278
Other post-secondary education	0.14	0.00	0.00	1.00	0.35	278
Panel B: Labor market background						
Any work in the 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)	44.95	0.00	0.00	702.75	102.07	277
Written contract	0.13	0.00	0.00	1.00	0.33	278
Panel C: Search behavior						
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	3.00	0.52	278
Panel D: Search alignment with comparative advantage						
Δ planned apps (w, aligned - misaligned)	0.74	2.00	-115.00	55.00	20.82	278
Δ % platform apps (aligned - misaligned)	-0.00	0.00	-1.00	1.00	0.21	278

Notes: **Table 1 shows baseline summary statistics for the tight field experiment.** Winsorized variables are winsorized at the 99^{th} percentile. All monetary values are in 2021 USD purchasing power parity terms.

rience of employment 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; Guvenen et al., 2020). Learning about skills by search is also difficult: only 3% jobseekers report ever receiving feedback about their skills during an unsuccessful job application.

Third, job search costs are high. One cause is the high cost of transportation in South Africa. Banerjee and Sequeira (2020) find that transport costs make up more than 50% of reported monthly income in a similar sample. Kerr (2017) documents average transport cost of 11% of income among a working population. In our sample, the average person spends half as much on search each week as the average employed person earns, covering transport, mobile phone data, and printing applications. This further limits the scope for jobseekers to learn about their skills through search and raises costs of misdirected job search.

2.2 Skill Assessments

Skill assessments in the tight experiment covered two domains: communication and numeracy. The numeracy assessment focuses on practical arithmetic and pattern recognition. It was developed by a large retail chain and used in their applicant screening process, as they believe it identifies some of the skills needed by cashiers. The communication assessment captures English language listening, reading and comprehension skills. The assessment was developed by a South African adult education provider (www.mediaworks.co.za) and is designed to assess English proficiency for high school students. Candidates also complete a measure of "concept formation" developed for the South African context (Taylor, 1994), similar to a Raven's test (Raven and Raven, 2003). This is a non-verbal measure of fluid intelligence, people's ability to see underlying commonalities across situations and to use logic in new situations. We describe the assessments, their psychometric properties, the assessment process, and the assessment result distributions in Appendix D. The distributions show a wide range of assessment scores and little evidence of ceiling or floor effects (Figure A1). These assessments are also used in the big experiment, along with three additional assessments that we describe in Section 4.2.

Based on the assessment results, we classify jobseekers into skill quintiles for both skills relative to a reference group. The reference group consists of roughly 12,000 similar jobseekers assessed by Harambee in Johannesburg using the same assessments. These jobseekers are young and completed high school in rural areas or townships around Johannesburg, just like our participants. The quintile scores are our main measure of jobseekers' skills throughout the analysis. We use ordinal measures of skill and skill beliefs throughout the paper. We prefer these to cardinal skill measures (e.g. 80%) because cardinal measures are dependent on assessment-specific scales and ordinal measures align better with our measures of skill beliefs and our definition of comparative advantage.[8]

[8]Reassuringly, communication and numeracy quintiles are positively associated with jobseekers' self-reported grades on their high school leaving exams in respectively English and mathematics (Table A1, columns 1-2). Unsurprisingly, exam grades do not perfectly predict quintiles of assessment scores, as the average jobseeker took the exam multiple years ago, the exam and assessments do not test identical skills, and jobseekers may not perfectly recall their

We classify a jobseeker as having a **comparative advantage in skill S** if they score in a higher quintile for the skill S assessment than the other skill. This captures the spirit of standard definitions of comparative advantage: for a given time endowment, each jobseeker has a comparative advantage in the skill they can supply more of to the market than other "suppliers:" other jobseekers from a similar background who are likely to be applying for the same jobs. This approach is relatively rare in labor economics because multidimensional skill measures are available in few datasets, but it follows existing work that uses similar data (e.g. Guvenen et al. 2020). This approach does not capture the jobseeker's rank in the skill distribution of applicants for any specific job, which would require very different data sourced from a large number of firms. Nor does it capture the jobseeker's relative rank in the skill distribution for the entire country, which would require nationally representative data and include many people with higher or lower levels of education who are not actually relevant competitors for the same jobs. But it does take advantage of the direct measures of multidimensional skills we have available and avoids the assumptions used in some other definitions of comparative advantage. For example, some labor economists estimate occupation-specific wages for either a single skill or for education level and use this to define comparative advantage in occupations based on levels of occupation or skill (e.g. Acemoglu and Autor 2011; Gibbons et al. 2005). These approaches have the advantage of pricing the value of different types of workers in different types of jobs, which we do not. But their wage estimates are conditional on the way workers currently sort into occupations. These are sensitive to a Lucas-style critique that wages might be different under a different type of sorting, which is precisely the type of mechanism we study.

We restrict our sample to jobseekers with a clear comparative advantage in skills, as noted earlier. This means that we drop the 24% of jobseekers who have the same quintile for both skills. This sample restriction is important because our experimental measures of skill-directed job search, described in the next section, can only be defined for jobseekers with a unique skill comparative advantage.

The scores in communication and numeracy are weakly related so these two assessments horizontally rather than vertically differentiate jobseekers: they show which jobseekers are better suited for specific types of jobs, rather than identifying jobseekers who are likely to be better at all or most job types. Table A2 shows that 62% of our sample has a comparative advantage in communication, 38% has a comparative advantage in numeracy, and only 13% are in the top two quintiles for both skills.[9]

exam scores. The difference between English and mathematics scores on the exam is also positively associated with jobseekers' comparative advantage in our assessments (column 3).

[9]Imposing the sample restriction means that communication and numeracy skill quintiles have zero correlation and hence reduces vertical differentiation in the sample. If this assumption is removed, communication and numeracy scores have correlation 0.31 and 17% of the jobseekers are in the top two quintiles for both skills. Communication and numeracy quintiles are both moderately correlated with fluid intelligence quintiles ($\rho \approx 0.25$ for both skills). The shares of jobseekers in this sample with communication and numeracy comparative advantage are not required to be equal because, as we note above, we define the quintiles using a larger sample of roughly 12,000 jobseekers assessed before the tight experiment.

2.3 Firms' Valuation of These Skills

Searching based on the skill comparative advantage we measure is most likely to help jobseekers under three conditions: firms value communication and numeracy skills; firms' relative valuation of the two skills varies, so it is optimal for different jobseekers to apply to different firms; and firms can at least partly observe jobseekers' skills. Here we present evidence consistent with each of these conditions to motivate the experimental analysis. This is not, of course, evidence that skill-directed job search improves labor market outcomes; the experiments are designed to do that.

Firms value communication and numeracy skills: Three patterns show that the skills we measure are valued by firms in this labor market, suggesting that jobseekers with skills more suitable to a vacancy may be more likely to be hired or offered higher wages.[10] First, client firms have paid Harambee to screen roughly 500,000 prospective workers using the communication, numeracy, and concept formation assessments, which we interpret as revealing a preference for using these skills in hiring. Second, Carranza et al. (2022) use multiple field experiments to show that firms are willing to pay to learn jobseekers' results on these skill assessments and that jobseekers who can certify their assessment results for communication, numeracy, and a few other skills have higher employment rates. Third, in an incentivized resume-ranking experiment, that we describe in detail in Appendix H, 91% of many firms preferred job applicants with higher communication or numeracy assessment results to job applicants with an additional six-month post-secondary training qualification.

Firms' relative valuation of skills varies: In the incentivized ranking experiment, 42% of the firms in our sample ranked the profile with high numeracy skills ahead of the profile with high communication skills and 58% 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' at least partly observe skills: To show that firms can partly observe skill levels, we conduct a measurement exercise embedded in a firm's hiring process. During the job search workshop used for the tight experiment, we ask jobseekers to choose between applying for a real communication or numeracy job at a firm that hires for a range of entry-level roles, including call center and data capture jobs. They prepare a resume 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 then evaluate every applicant for both jobs. Evaluators are blind to which applicant applied for which job and are not shown applicants' skill assessment results. We receive data on the evaluators' assessment of each applicant's communication skills, numeracy skills, and suitability for each type of job, as well as whether they recommend interviewing each candidate for each job.

[10] We do not claim that these are the only or necessarily even the most important skills in this labor market. Employers may also value the noncognitive skills we introduce in Section 4, social skills, physical strength, etc. We simply argue that communication and numeracy skills are, on average, valued by firms in this labor market.

Table 2: Employer Evaluation of Job Applicants Based on Skills

	Mean			Difference			Obs.
	Aligned	Non-aligned	SD	Δ	Δ/SD	p($\Delta = 0$)	
Panel A: Skill levels							
Skill (1-5)	2.93	2.79	0.66	0.15	0.22	0.01	277
Panel B: Job-related evaluation							
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 2 shows that an employer evaluates applicants more highly if their assessed comparative advantage in skills matches the job's requirements.** It compares employers' evaluation of jobseekers for two jobs: one that aligns with their comparative advantage in skills and one that does not. These evaluations are based on real applications that jobseekers in the tight experiment submit to a real employer listing both a communication-heavy job and a numeracy-heavy job. We ask the employer to evaluate applicants for both positions without knowing which position applicants applied to.

This measurement exercise shows that the firm's HR team can identify levels of jobseekers' skills. Table 2, panel A, column 6 shows that their assessments of skills are positively but not perfectly correlated with our measures of skill: they assign a 0.22 standard deviation higher score to the skill we assess as higher ($p = 0.01$). Their 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 are 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 6). This is a 26% increase relative to a 34% 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, suggesting the possibility of greater observability in natural job search where jobseekers can tailor their applications.

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

2.4 Beliefs about Skills and Comparative Advantage

We measure jobseekers' pre-treatment beliefs about their communication and numeracy skill quintiles in two ways: **general skill beliefs** and **beliefs about assessment results**. General skill be-

liefs are measured before taking the assessments. We ask jobseekers which quintile their general communication and numeracy skills are in, relative to the reference group of jobseekers from Johannesburg aged 18-34 with a matric (the high school leaving exam) from a township or rural school.[11] We define communication skills and numeracy skills before asking the question.[12] Beliefs about assessment results are measured after taking the assessments. We ask jobseekers which quintile they scored on the communication and numeracy assessments, relative to the same reference group. The questions about general skills and assessment results capture two related but not identical concepts, as jobseekers may, for example, believe that they have high numeracy skills but performed poorly in this specific assessment. We focus on beliefs about general skills for most of the paper, except where we specifically note otherwise. In practice, these two beliefs are strongly positively correlated within skill, suggesting that jobseekers view the assessments as relevant to their general skills.[13]

We summarize the multidimensional relationship between assessed and believed skills in two statistics. First, **fraction aligned beliefs**: the average of two indicators, one for communication and one for numeracy, each equal to one if perceived and assessed skill quintiles are equal. Second, **aligned comparative advantage belief**: an indicator equal to one if the jobseeker has a higher perceived quintile for the skill in which they have a higher assessed quintile. These two beliefs are positively correlated by construction but have some separate variation, generating a correlation coefficient of 0.21. To understand the separate variation in the two measures, consider a candidate whose perceived and assessed skills are in quintiles 2/4 for communication and numeracy. She will score 1 on both of the two belief measures. Raising her numeracy without changing her communication beliefs will decrease the fraction of aligned beliefs to 0.5 without changing aligned comparative advantage, while just lowering her numeracy belief will decrease the fraction of aligned beliefs to 0.5 and eventually flip aligned comparative advantage belief from one to zero. The joint distribution of assessment results and beliefs across two skills is a complex ob-

[11] The specific wording is: "Think about 100 people who are job-seekers from Johannesburg aged 18-34 with a matric from a township or rural school. Imagine that we rank everyone according to their numeracy skills. We create five equal size groups. The first group are the 20 people with the highest numeracy 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 lower 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?" Jobseekers receive a visual aid, a drawing of five quintiles similar to the report used to provide information on jobseekers' skills.

[12] The specific wording is: "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."

[13] Regressing general skill belief on belief about assessment results produces coefficients of 0.39-0.52 across the two skills, with or without controls for assessment results and demographic characteristics. General skill beliefs about communication and numeracy are also positively correlated with jobseekers' self-reports about their results on the high school graduation exams in respectively English and mathematics (Table A1, columns 4-5). Beliefs about comparative advantage are positively associated with the difference in scores between the two exam subjects (columns 6-7). However, exam scores do not perfectly predict beliefs, showing that jobseekers' beliefs depend on multiple sources of information.

Table 3: Baseline Beliefs About Skills

	(1) Mean	(2) SD	(3) Obs.
Panel A: Aligned comp. adv belief			
General skill	0.49	0.50	278
Assessment result	0.54	0.50	278
Panel B: Fraction aligned beliefs			
General skill	0.22	0.31	278
Assessment result	0.17	0.25	278

Notes: **Table 3 shows a large degree of misalignment between baseline beliefs about skills and skill assessment results.** General skill refers to jobseekers' beliefs about their general communication and general numeracy skills. Assessment result refers to jobseekers' beliefs about their performance on the assessments administered as part of the experiment. The reference population are a group of roughly 12,000 jobseekers from Johannesburg aged 18-34 with a matric from a township or rural school. Panel A shows an indicator equal to one if the jobseeker has a higher perceived quintile for the skill in which they have a higher assessed quintile. Panel B shows the average of two indicators, one for communication and one for numeracy, each equal to one if perceived and assessed skill quintiles are equal.

ject that can be summarized in many different ways. In Section 5.1, we also consider how much jobseekers' believed skills are higher or lower than their assessed skills, although this is not the primary focus of this paper. In Appendix A.2 we consider continuous alternatives to the fraction aligned beliefs and aligned comparative advantage belief.

We use these beliefs measures to highlight four key empirical patterns that motivate our conceptual framework and subsequent analysis.

Jobseekers' believed and assessed comparative advantage often differ. Before taking assessments, only 49% of jobseekers' beliefs about their comparative advantage in general skills align with their comparative advantage on our skill assessments (Table 3, panel A). This is slightly better than the 40% of aligned beliefs that would arise from random guessing of skill quintiles.[14] Jobseekers' beliefs about their results on our specific assessments are similarly misaligned: only 54% percent of jobseekers assess their relative assessment performance correctly. Thus even taking the assessment and knowing the type of content associated with each skill is not sufficient to align jobseekers' beliefs about their comparative advantage with their assessed skills.

Jobseekers also have misaligned beliefs about their skill levels (Table 3, panel B). Averaging

[14]The random benchmark is below 50% because our sample has by definition no ties in their skill quintiles. This implies that jobseekers with tied beliefs are always coded as misaligned.

across all jobseekers, only 22% of believed general skill quintiles equal assessed skill quintiles and only 17% of believed assessment result quintiles equal assessed skill quintiles. This pattern is driven more by overconfidence than underconfidence: 60% of jobseekers' perceived general skills are above their assessed skills and only 17% are below, in line with patterns documented in many other settings (Santos-Pinto and de la Rosa, 2020).

This pattern motivates a conceptual framework where jobseekers' skills and beliefs about their skills may differ, potentially leading their comparative advantage and perceived comparative advantage to differ.

We observe limited gender heterogeneity in the relationship between measured and perceived skills in both the tight and big experiments (Tables A3 and A4). Conditional on demographic differences between men and women in our sample, the observed gender differences are small in both experiments and marginally statistically significant in only the big experiment. There are no differences in the fraction of over- or underconfident beliefs, matching the pattern in a recent metastudy of gender and confidence (Bandiera et al., 2022). Hence, we focus on gender-pooled results throughout this section.

Differences between believed and measured skills are persistent: To analyze the persistence of misaligned beliefs over time, we use the big experiment's follow-up survey, which occurred on average 3.5 month after baseline. The share of control group jobseekers whose believed and assessed comparative advantage align changes little from baseline to the endline, even for those with above-median search effort and those who are employed (Table A5, columns 1-3). The same persistence holds for the fraction of aligned beliefs (columns 4-6).

The strong persistence of misaligned beliefs emphasizes that this labor market offers limited scope for these types of jobseekers to learn about their skills, at least for the time horizon we study. It suggests that a dynamic conceptual framework incorporating belief updating through search or work is not needed in this setting.

Jobseekers believe skills are valued: Two patterns in survey data from big experiment show that jobseekers believe there are positive labor market returns to skills. First, individual jobseekers think their outcomes would improve if their skills were higher. We asked each jobseeker their expected search duration and earnings conditional on finding a job and then asked their expectation for another jobseeker who had better numeracy skills but was otherwise identical to themselves. Control group 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 118.46 USD PPP more (13% of the mean, $p < 0.01$). We did not ask these questions about communication skill due to time constraints.

Second, jobseekers with higher skill beliefs expect to earn more and get a job faster. This is a between-jobseeker comparison, so it is vulnerable to omitted variable bias, but the pattern is robust to controlling for actual skill levels, age, gender, education, and work history. A one standard deviation higher average skill belief is associated with 5% higher expected wage ($p < 0.001$) and 3.5% shorter expected search duration ($p < 0.001$) (Table A6).

These two patterns motivate a conceptual framework where information about skills might alter beliefs about the probability of getting a job and wages conditional on getting a job.

Beliefs about comparative advantage predict how search is directed: Finally, we show that jobseekers' skill beliefs are correlated with their job search direction. In the field experiment, we ask candidates what skill is most valuable for the types of jobs they are applying for. In the control group, 47% and 22% answer communication and numeracy respectively and these answers are strongly linked to their beliefs about their comparative advantage (Table A7). Jobseekers are 6-10 percentage points more likely to state that the skill of their comparative advantage was most important for the jobs that they are applying for compared to choosing any of the other skills (17-45% of the mean, $p < 0.01$. The correlations are robust across skill domains and to controlling for measured comparative advantage and demographic controls. We find a similar pattern in the tight experiment using different measures of search direction, which we introduce in Section 3.4.

This link between search direction and comparative advantage beliefs motivates an analytical framework that studies how jobseekers' perceived comparative advantage in skills shapes the direction of their search effort between jobs with different skill requirements.

2.5 Conceptual Framework

Our conceptual framework is directly designed to show how the skill belief patterns we have documented can influence job search activities and outcomes, in a similar spirit to Mahoney (2022). In this framework, jobseekers have multidimensional skills, have potentially inaccurate beliefs about their skill levels and comparative advantage, make choices about how to direct search effort across different job types, and receive job offers with specific wages. We assume here that total search effort is fixed so we can focus on the search direction analysis but relax this assumption in Section 5.1. Given the limited evidence of learning about skills we document above, we use a static framework and do not model dynamic belief updating in response to search outcomes. Given the types of data we observe, we use a partial equilibrium framework focusing on jobseekers' search and beliefs but treating labor demand and wage posting as fixed. This framework has a similar spirit to recent models of partially directed job search, where jobseekers try to direct search to highest-wage vacancies but face uncertainty about wages, but is more stylized and not designed for estimation (Lentz et al., 2022; Wu, 2021).

In this framework, each jobseeker is endowed with communication skill level S_C and numeracy skill level S_N.[15] Each job demands primarily communication or primarily numeracy skills. Jobseekers split fixed total search effort $\bar{E}$ between search for communication jobs E_C and numeracy jobs E_N. Given skills and search effort, jobseekers receive an offer for job type j with probability $P_j(S_C, S_N, E_j)$ with associated "wage" $W_j(S_C, S_N, E_j)$, which is a reduced-form expression for the expected present value of the offer. We define $V_j(S_C, S_N, E_j) = P_j(S_C, S_N, E_j) \times W_j(S_C, S_N, E_j)$

[15] We assume that the two skills have the same distributions, so absolute and comparative advantage are aligned. This simplifies the algebra and matches our data, which measures skills in both domain using rankings not levels.

and assume V_j is increasing and concave in all three arguments and that

$$\partial V_j/\partial S_j > \partial V_j/\partial S_i > 0. \tag{1}$$

for $j \neq i$. This assumption allows both types of jobs to value both skills, but each type of job to value one skill more than the other. We also assume that skill and search effort are technical complements and are 'more complementary' within than across dimensions. Intuitively, this means that a jobseeker with high numeracy skills will get a higher return to directing marginal search effort to numeracy jobs than communication jobs and vice versa. Formally,

$$\frac{\partial^2 V_j}{\partial S_j \partial E_j} > \frac{\partial^2 V_i}{\partial S_j \partial E_i} > 0 \tag{2}$$

for $j \neq i$. Finally, we assume that the gross utility from job search $U(V_C, V_N)$ is increasing and concave in both arguments. We simply assume that jobseekers derive utility from offers, rather than modeling the offer acceptance decision or reservation wage.

The jobseeker directs search effort to equalize the marginal utility of searching for each job type:

$$\frac{\partial U}{\partial V_C} \times \frac{\partial V_C}{\partial E_C} = \frac{\partial U}{\partial V_N} \times \frac{\partial V_N}{\partial E_N}, \tag{3}$$

where $\frac{\partial U}{\partial V_j}$ captures the jobseeker's preferences over non-pecuniary aspects of the job types. 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 assumptions (1) and (2), the first term is more steeply increasing in communication skill than the second term. So, to satisfy the first order condition in (3), a jobseeker with higher communication skills, i.e., with comparative advantage in communication will direct more effort to searching for communication than for numeracy jobs.

We assume each jobseeker has beliefs about their skill levels $\tilde{S}_C$ and $\tilde{S}_N$. In this framework, jobseekers direct search based on their skill beliefs but their search outcomes depend on their actual skills. Or, in a slightly more general framework, search outcomes could depend on noisy signals of skills received by firms. This framework delivers two testable predictions that we evaluate in our experiments:

P1 Jobseekers will direct more search effort to jobs requiring the skill in which they have a perceived comparative advantage to satisfy condition (3), irrespective of their actual comparative advantage.

P2 Jobseekers whose perceived comparative advantage does not match their actual comparative advantage will receive fewer offers, lower wages conditional on offers, or both, due to assumption (2).

We use multiple measures of search effort and outcomes, defined below, to test these predictions.

3 Tight Experiment: Effects on Beliefs and Behavioral Measures of Directed Search

In this section, we report the results from the tight experiment, which is designed to precisely measure skills-directed job search with actual job choices and to collect rich beliefs data. We show that giving jobseekers their assessment results shifts their believed comparative advantage toward their assessed comparative advantage and redirects their search effort toward jobs whose skill requirements align with their assessed comparative advantage.

This experiment is optimized to collect detailed and comprehensive measures of search direction, at the jobseeker×job level. This is difficult in labor market research that relies on survey data from jobseekers, which can measure self-reports on general search strategies but struggle to capture precise measures of applications to specific jobs.

3.1 Experimental Design

As discussed in Section 2, we work with a sample of young jobseekers who have signed up to receive job assistance services from a social enterprise. On the day of the workshop, jobseekers respond to a pre-treatment survey, take assessments of their numeracy, communication, and concept formation skills and complete a brief survey about their perceived performance on the assessments.

To cleanly define beliefs about comparative advantage and job choices aligned with comparative advantage in skills we focus our analysis on the 278 individuals with a clear comparative advantage in one of the skills. Jobseekers who have the same quintile in both skills have no clear comparative advantage, so we cannot sensibly define their job search choices as aligned or misaligned with their comparative advantage.[16] After restricting our sample to jobseekers with a comparative advantage in skills, we assign 139 jobseekers assessed over 17 days to the treatment group and 139 jobseekers assessed over a different 17 days to the control group. We randomize treatment by assessment date to reduce risks of spillovers between treated and control jobseekers. Treatment assignments are balanced on baseline covariates, both in the full sample and in the subsample with a clear comparative advantage in skills (Table A10).

Treated jobseekers receive 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 watch a video that explains the content of the report and how to interpret the quintiles. They can ask questions during and after the video. Jobseekers in the control group do not receive a report. They watch the part of the treatment video explaining the assessments, to hold constant any effect of having better information about communication or numeracy skills on beliefs or search behavior. But they do not watch the part of the treatment video explaining the report content.

The intervention is designed to give jobseekers information about their skills relative to other

[16] However, including these jobseekers in the sample and classifying all their job search choices as misaligned with comparative advantage produces qualitatively similar treatment effects on skill beliefs (Table A8) and search (Table A9).

Figure 1: Sample Report

REPORT ON CANDIDATE COMPETENCIES
-Personal Copy-

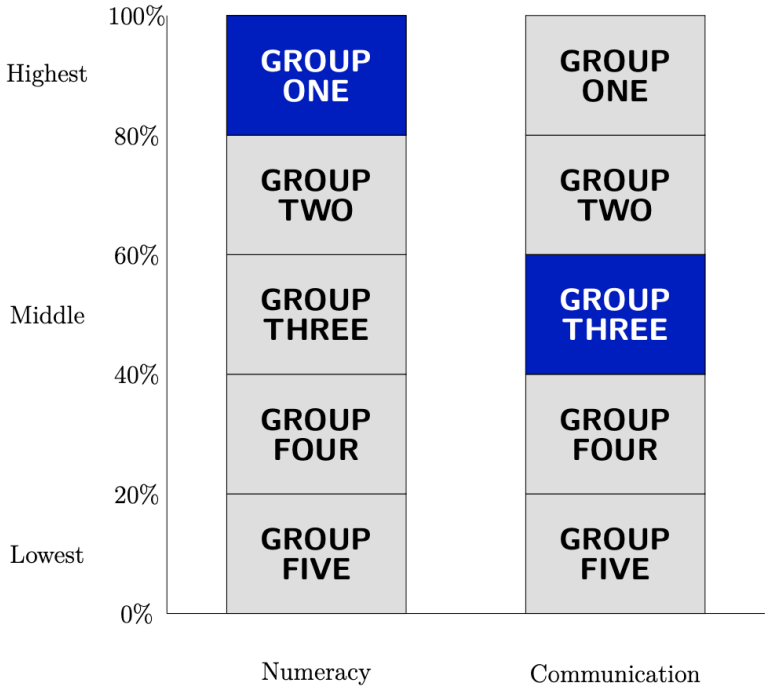
This report contains results from the assessments you took today. These results can help you learn about some of your strengths and weaknesses and inform your job search.

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 and no branding.

jobseekers. It encourages them to think about what jobs will value their skills but it does not encourage them to apply to any specific jobs. Recognizing that processing assessment results can be difficult, we iteratively piloted multiple versions of the reports and explanatory videos. After the video, we ask treated jobseekers a comprehension check: to report what the assessment results show about their skill comparative advantage. 95% of treated jobseekers answer this comprehension check correctly. Taking the assessments and answering skill-focused survey questions might prime jobseekers to think more about the role of skills in job search during the experiment. This might increase the weight they place on skill levels and skill match during the post-treatment surveys and tasks. However, this effect should be largely similar between treatment and control groups.

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 other identifying information. It is not branded by Harambee or any of the research institutions. It also provides no information about the relationship between the assessments and productivity and no information about firms' existing use of these assessments. We show in Section 5.4 that information acquisition by firms is unlikely to explain the results of the experiment.

3.2 Specification

We estimate treatment effects of receiving information about one's relative ranking on skills at the individual level. For these outcomes we estimate:

$$Y_{id} = T_d \cdot \beta + \mathbf{X}_{id} \cdot \Gamma + \epsilon_{id}, \tag{4}$$

where the average treatment effect, β 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.[17] We use heteroskedasticity-robust standard errors clustered by assessment date, the unit of treatment assignment. We show in Table A11 that our hypothesis test results are robust to using a wild cluster bootstrap to account for the low number of clusters (34) and to adjusting for multiple hypothesis testing using sharpened q-values, following Benjamini et al. (2006).

We also test if increased alignment between perceived and measured skills and increased alignment of search direction by skill is concentrated among jobseekers whose baseline beliefs are not aligned with their comparative advantage. To show this, we estimate models of the form:

$$Y_{id} = T_d \cdot \alpha^{misaligned} + T_d \cdot Aligned_{id} \cdot \alpha^{aligned} + Align_{id} \cdot \delta + \mathbf{X}_{id} \cdot \Gamma + \epsilon_{id}, \tag{5}$$

where $Aligned_{id}$ is an indicator for jobseekers whose beliefs about their comparative advantage

[17] $\mathbf{X}_{id}$ contains age; a dummy for being female; dummies for having a high school leaving certificate, 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; dummies for having a comparative advantage in either of the two skills; a pre-treatment value of the outcome Y_{id} where available; and block fixed effects, to account for the fact that we randomly assign days to treatment groups within blocks of 4 sequential days.

on the assessments match their measured comparative advantage at baseline.[18] These jobseekers receive less information from treatment. The main objects of interest are the average treatment effect for jobseekers with nonaligned baseline comparative advantage beliefs, $\alpha^{misaligned}$, and the difference in average treatment effect between jobseekers with misaligned and jobseekers with aligned baseline beliefs, $\alpha^{aligned}$.

As in any heterogeneity analysis, $Aligned_{id}$ may be correlated with other jobseeker-level characteristics, complicating the interpretation of $(\alpha^{misaligned}, \alpha^{aligned})$. But this is a relatively minor concern in this specific application. We show in Table A12 that having an aligned comparative advantage belief at baseline is unrelated to gender, age, employment, and work experience: these variables jointly explain only 3% of the variation in aligned comparative advantage beliefs. This belief is, unsurprisingly, related to the communication and numeracy assessment scores, but these scores explain only 15% of the belief variation. Our key treatment effect estimates are robust to using a non-binary measure of aligned comparative advantage beliefs (Table A13). In Section 5.1, we explore heterogeneous treatment effects by the combination of baseline-aligned comparative advantage belief and over-/underconfidence in skill level, showing that this analysis does not substantially change the findings of the paper.

Both estimating equations, all baseline covariates, and most outcome measures are prespecified at https://doi.org/10.1257/rct.10000-1.0. We describe the relationship between the preanalysis plan and our final analysis in Appendix F.

3.3 Information About Skills Aligns Beliefs with Comparative Advantage

Jobseekers who receive information about their skills are more likely to have beliefs about their skills that are aligned with their assessed comparative advantage. Table 4 presents results for general beliefs about skills. Treated jobseekers are on average 13.5pp more likely to report beliefs about their comparative advantage which are aligned with their measured comparative advantage, a 28% increase from the control mean of 47.5% (column 1, $p = 0.001$).[19]

Most of the treatment effect on comparative advantage belief alignment is driven by jobseekers whose beliefs were not aligned at baseline (column 2). The coefficient on the treatment dummy shows the effect on those whose beliefs were not aligned at baseline: in column 2 of Table 4, the likelihood that initially misaligned participants' comparative advantage beliefs will align with their measured comparative advantage increases by 21pp (a 44% increase over the control mean, $p < 0.001$). The treatment effect for those whose beliefs were aligned at baseline – which is the sum of the coefficients on the treatment indicator and the interaction term – is given in the row below the main coefficients, an insignificant 7pp. The difference between jobseekers with aligned and non-aligned baseline beliefs is a marginally insignificant 14pp, or 30% of the baseline con-

[18] We choose to use beliefs about assessment results for heterogeneity because they indicate who will receive a signal contradicting prevailing beliefs. This is the only place in this section that we use assessment-specific beliefs rather than general skill beliefs. Estimating heterogeneity using the measure of beliefs about general skills yields very similar results.

[19] The elicitation of the beliefs is described in Section 2.2.

Table 4: Treatment Effects on Beliefs About Skills - Tight Experiment

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

Notes: **Table 4 shows that the treatment aligns beliefs about skills with their assessed skills in the tight experiment.** Columns 1 and 2 show effects on a dummy indicating beliefs about respondents' comparative advantage in skills that are aligned with the assessment results. Columns 3 and 4 show treatment effects on the fraction of skill quintile beliefs that are aligned with the assessment results. Columns 2 and 4 show treatment effect heterogeneity by whether jobseekers had aligned comparative advantage beliefs at baseline. Controls include randomization block fixed effects, and prespecified baseline covariates (skill quintile dummies, age, gender, and dummies for having completed a university degree, other post-secondary education, and high school). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

trol mean, given by the "Treatment × Aligned comp adv belief" coefficient. Our model predicts that increased accuracy about their comparative advantage in skills should have consequences for jobseekers' search behavior and labor market outcomes.

Treatment shifts the level of skill beliefs as well as beliefs about skill comparative advantage. Treated jobseekers' believed skill quintiles are 7.8 percentage points more likely to match their assessed skill quintiles, a 43% increase on from the control group mean (Table 4, column 3, p=0.005). Heterogeneity in treatment effects by baseline comparative advantage beliefs goes in the opposite direction, jobseekers with initially aligned comparative advantage beliefs increase the fraction of aligned beliefs by 8 percentage points more than the control group. Importantly, there are no strict theoretical reasons to expect such heterogeneity to align, as aligned comparative advantage beliefs and aligned skill level beliefs can and do differ for many jobseekers.

In addition to these main results, we report several robustness checks and additional results

that may interest the reader but are not the primary focus of this paper. Treatment effect estimates are robust to using more continuous measures of skill beliefs (Table A13) and hypothesis test results are robust to wild cluster bootstrap and multiple hypothesis testing adjustment (Table A11). Heterogeneous treatment effects are robust to controlling for baseline 'confidence' – the difference between belief and measured skill quintiles – and its interaction with treatment (Table A14). Treatment reduces the variance of beliefs about skills, constructed by measuring jobseekers' belief distribution over all the quintiles, consistent with jobseekers acquiring more precise information from treatment (Table A15 and Figure A2). Belief updating reflects beliefs about both absolute performance on the assessments and performance relative to the reference group, suggesting that jobseekers learn about both their own skills and their place in the skill distribution (Table A16). Treatment updates underconfident beliefs more than overconfident beliefs (Table A17), matching asymmetric belief updating patterns in other research (e.g. Zimmermann 2020). Belief updating does not differ by jobseekers' cognitive skills, proxied by the concept formation assessment score (Table A18), speaking to a literature on the role of cognitive skills in information processing and belief formation (e.g. Kahan (2013); Kruger and Dunning (1999)).

Finally, we find no evidence of differential belief updating by gender (Table A19). This, combined with the lack of gender heterogeneity in baseline skill beliefs, means that we pool both genders for the rest of the paper.

3.4 Job Search with Better Aligned Beliefs about Skills

We now use this exogenous variation in beliefs to study how misaligned skills and skill beliefs influence job-search behavior. Guided by our conceptual framework, we interpret the treatment effects on skill beliefs as the first stage in a causal chain from treatment to skill beliefs to search activities and outcomes. However, we do not regress search on beliefs instrumented by treatment because the exclusion restriction holds only under the strong assumption that there is a single measure of beliefs that captures all updating. This follows standard practice in beliefs experiments, reviewed by Haaland et al. (2023).

We find robust evidence that jobseekers redirect search toward jobs that align with their assessed comparative advantage across four measures of skill-directed search. These differ on several features, such as realism, incentivization, timing, and susceptibility to experimenter demand effects.

Job choice task: We design a novel incentive-compatible job search task, in which jobseekers choose which realistic job advertisements with different skill requirements they would apply to (the skill requirements of job advertisements are coded by recruiters). In this task, we asked respondents to make 11 application choices between pairs of job advertisements. After the treatment giving participants their assessment results (or the control video giving them information about the assessments), participants were shown each pair of advertisements on a piece of paper, given time to read them, and asked by the enumerator to select which they would like to apply to.

The job advertisements are based on real job advertisements on the job search platform run by

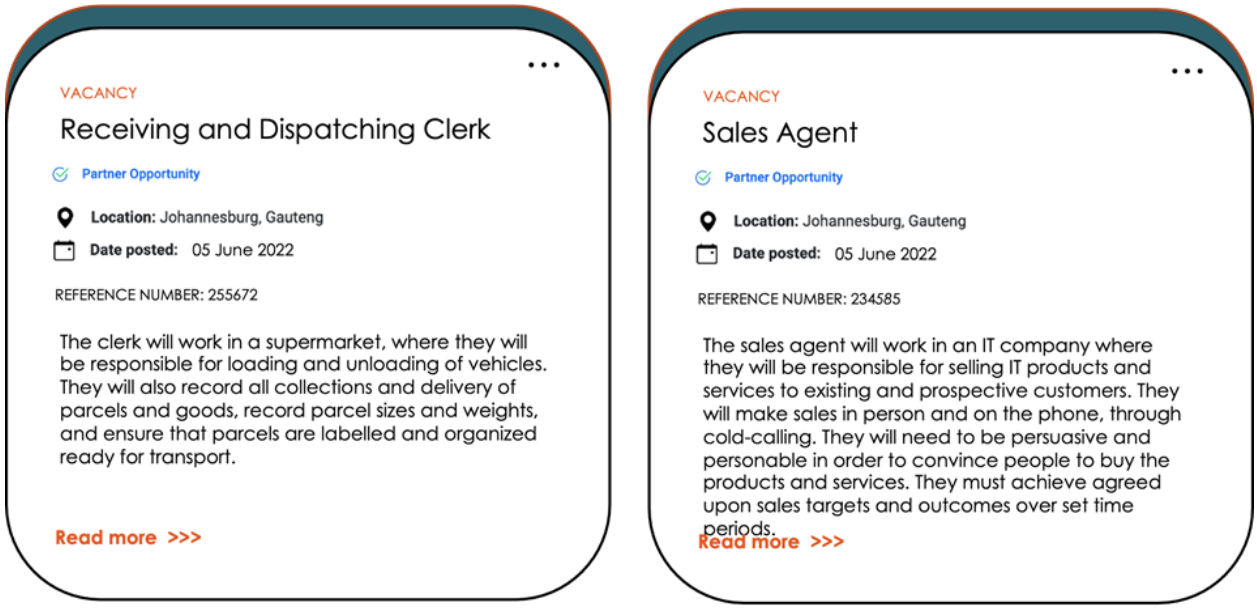


Figure 2: Sample Pair of Jobs from Job Choice Task

our implementation partner (SAYouth.mobi). We started from real advertisements, standardized their format, and removed information on the employer name and location to reduce their role in jobseekers' job choices. We used only job advertisements for entry-level positions that required no special education or qualifications (such as a driving license) or ability to use machines.

All jobseekers were platform users so the choice represented a real-life choice they often made during job search. The average jobseeker visited the platform on 3 days of the 30 days before their workshop and applied to 6.5 jobs (Table A20).

Importantly, each pair consisted of a job that required primarily numeracy skills and a job that required primarily communication skills. We focus on jobs that have a clear skill requirement, recognizing that this is not the case for all jobs to which jobseekers might apply. We selected a long-list of 28 job adverts from the platform that we thought might require one or other skill. To classify jobs into numeracy and communication jobs, we asked 13 human resources and recruitment staff with experience hiring for entry-level roles to classify jobs according to their communication and numeracy skill requirements, expected wage, overall and gender-specific desirability, and transparency of skill requirements. We then created pairs of jobs with opposite skill requirements and similar desirability levels. Figure 2 shows an example pair and Table 5 shows the full list of job titles. Participants viewed the 11 job advertisement pairs in a randomized order.

The pairing exercise produced jobs that differ on skill requirements but are similar in other respects, allowing us to isolate the role of skill-directed job search. Averaging across all pairs, our 13 placement officers scored the job we defined as numeracy-heavy as needing 2.7 standard deviations more numeracy skills and likely to pay only 0.3 standard deviations more.

We incentivize jobseekers to respond truthfully to the job choice task in two ways. First, one pair of the 11 job advertisements were live advertisements for jobs with a partner firm. Jobseekers

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

Table 5: Job titles used in the job choice task

were told we would submit an application on their behalf to their favored advertisement from this pair of jobs after the workshop and that they would receive assistance in writing their CV and cover letter for the job at the end of the workshop. Jobseekers did not know which pair of jobs was live. All other pairs of jobs were taken from real advertisements but we did not forward applications to these jobs, partly because some had deadlines before the experiment finished. We further incentivized jobseekers' decisions by telling them that at the end of the workshop, we would give them a list of entry-level job titles similar to the jobs they preferred in the task to assist them in job search in the future. Jobseekers received the list of job titles at the end of the workshop as promised.

Importantly, the treatment gives information on only two skills, allowing a simple measurement of skill-directed job search. The big experiment described in the next section uses more skills, which creates an arguably more realistic environment but makes it more complex to define skill-directed job search. The tight and big experiments are thus designed to provide respectively clearer evidence on a mechanism and more realistic evidence on outcomes.

We measure the impact on the share of aligned applications: the share of total job choices that were in line with participants' measured comparative advantage. Column 3 of Table 6 shows that treated jobseekers have, on average, a 4 percentage points higher share of job choices that align with their measured comparative advantage ($p = 0.29$), compared to the control mean of 55% of jobseekers making an aligned job choice. This effect hides important heterogeneity described below.

We leverage the fact that we use only two skills in the tight experiment to cleanly define a baseline measure of aligned comparative advantage beliefs and estimate whether the effects on search direction differ by whether jobseekers had aligned beliefs at baseline. Respondents with aligned comparative advantage beliefs at baseline should not change their search direction in response to the treatment, while those with unaligned beliefs should change their search direction. This is indeed the case. Individuals with initially unaligned beliefs have 8.8 percentage points higher

share aligned job choices ($p = 0.024$), a 16% increase relative to the control mean. Individuals with initially aligned beliefs have a significantly lower treatment effect with a point estimate of -2 percentage points (Table 6, column 2). This heterogeneity is consistent with the heterogeneity in comparative advantage beliefs: jobseekers with initially unaligned beliefs about their comparative advantage in skills were much more likely to update their beliefs to align with their comparative advantage than those whose beliefs were initially aligned. The observed treatment effects are also in line with the model predictions of a shift in search effort towards jobs that align with jobseekers' comparative advantage, particularly among individuals with initially unaligned beliefs.

The incentivized job choice task is highly controlled and it is difficult for any other behavioral mechanism to be causing the observed treatment effects. In the job choice task, the job pairs are carefully selected so that each pair is similar in other characteristics, including expected wages, as we explain on page 25. This makes it likely that the change in job choices is driven by jobseekers being more likely to apply to jobs aligned with their skills. In a field setting, in contrast, it is possible that jobs with different skill content also differ on unobserved characteristics which drive observed search direction. Similar challenges apply to work studying how job applications on search platforms respond to posted wages (Belot et al., 2022b; He et al., 2021). The task is also incentive-compatible, providing some protection against experimenter demand effects.

The search direction results so far suggest that jobseekers have some information about the relative skill demands of different jobs. However, they might still face information frictions about the jobs' skill requirements. We provide an additional test by explicitly revealing relative skill demand for the last two pairs of jobs (recall that order is randomized). Table A21 shows that jobseekers with initially misaligned beliefs align their search for job choices with and without revealed skill requirements, but substantially more when skill requirements are revealed. This suggests that jobseekers have some but not incomplete information about the skills requirements of some of the jobs.[20]

Clicks on links to real jobs: To study the effects on search direction outside the workshop, we send jobseekers text messages with links to real job opportunities on SAYouth.mobi about a week after the workshop. We send each jobseeker three text messages in random order, one numeracy job, one communication job, and one job that aligned with the skill content of the majority of their choices in the job-choice task. The messages only contain the job title and do not contain any further information about the skill requirements of the advertised jobs.[21] We track whether jobseekers clicked on these links. We use this data to construct a second measure of directed search: the difference in the click rates for jobs that are aligned with respondents' comparative advantage and jobs that are not aligned with their comparative advantage. An increase in this variable means that jobseekers exert more search effort toward jobs aligned with their comparative

[20] In Appendix C we describe an additional part of the experiment that shows that there are no treatment effect on jobseekers' willingness-to-pay for information about vacancies' skill demand.

[21] The message text was: Hi [firstname], we are contacting you in relation to the workshop you participated in at the [venue of workshop]. We found a job opportunity for a [job title] post on the SAyouth.mobi site that you might be interested in. The job is within traveling distance of [area around venue of workshop]. Please follow the link if you would like to apply.

Table 6: Treatment Effects on Search Direction - Tight Experiment

	Aligned search index		% aligned (job choice)		Δ SMS click rate		Δ planned apps (w)		Δ % aligned platform apps	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment	0.269**	0.603***	0.037	0.088**	0.071	0.157	1.420	4.746**	0.063**	0.089*
	(0.103)	(0.140)	(0.034)	(0.038)	(0.061)	(0.096)	(1.261)	(1.763)	(0.023)	(0.048)
Treatment × Aligned comp adv belief (bl)		-0.648***		-0.104**		-0.163		-6.629**		-0.048
		(0.205)		(0.039)		(0.128)		(2.651)		(0.080)
Aligned comp adv belief (bl)		0.727***		0.165***		0.100		8.697***		0.013
		(0.153)		(0.035)		(0.102)		(2.412)		(0.049)
Treatment effect:		-0.045		-0.016		-0.005		-1.883		0.041
Aligned comp adv belief (bl)		(0.131)		(0.034)		(0.081)		(1.807)		(0.044)
Control mean	-0.000	-0.000	0.550	0.550	-0.032	-0.032	4.331	4.331	0.006	0.006
Observations	278	278	278	278	278	278	278	278	278	278

Notes: **Table 6 shows that providing jobseekers with information about their relative comparative advantage in skills aligns their search direction with their assessed comparative advantage in skills in the tight experiment.** Aligned job search is defined as directing search effort toward jobs that mostly require the skill that aligns with job seekers' measured comparative advantage. Columns 1 and 2 show impacts on an index of search direction. The index is constructed as the variance-covariance weighted average of the search alignment measures displayed in columns 3 to 10 (Anderson, 2008). Columns 3 and 4 show the impact on the percentage of 11 incentivized job choices that are aligned with the comparative advantage. Columns 5 and 6 show the effects on the difference in link click rates between aligned and non-aligned jobs sent to job seekers via SMS. All job seekers receive three links with at least one job for each skill. Columns 7 and 8 show effects on the difference between aligned and non-aligned applications (winsorized). Columns 9 and 10 show effects on the difference in between the percentage of aligned and non-aligned applications on the online job search platform SAyouth.mobi. Jobs are classified using a list of common numeracy-heavy and communication-heavy common jobs classified by experts. Even columns show heterogeneity by whether individuals have aligned comparative advantage beliefs at baseline. Controls include randomization block fixed effects, and pre-specified baseline covariates (skill quintile dummies, age, gender, dummies for having completed a university degree, other post-secondary certificate or diploma, and high school, and a dummy indicating a comparative advantage in numeracy). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

advantage in skills.

We again find an increase in skills-based search direction. Column 3 of Table 6 shows that the treatment shifts the difference in click rates by 7pp towards opportunities with comparative advantage ($p = 0.25$). We again observe that the treatment effect is fully driven by jobseekers with initially unaligned beliefs. They exhibit a treatment effect of 16 percentage points ($p = 0.11$) while individuals with initially aligned beliefs have a point estimate of zero. However, the results are less precisely estimated because our estimation uses one observation per participant, compared to 11 observations per participant in the job choice task. While not conclusive, these results suggest that jobseekers adjust their search effort even when just seeing job titles to make that decision. This measure captures whether jobseekers are able to target search according to the skill requirements of jobs based on very limited information, namely the job title.

Planned applications for numeracy and communication jobs: We survey participants about the number of jobs they plan to send to communication-heavy and numeracy heavy in the 30 days after the workshop. We use these data to construct the difference in planned aligned applications vs non-aligned applications, winsorized at the 99^{th} percentile. Treated individuals have a 1.42 larger difference in the number of planned aligned applications vs non-aligned applications (Table 6, column 7, $p = 0.27$). The effect is stronger for people with initially misaligned comparative advantage beliefs, who see a significant 4.75 increase in the difference between planned aligned and non-aligned applications (column 8, $p = 0.011$). The observed change in search direction is a result of both a slight increase in planned aligned applications and a reduction in planned non-aligned applications. This measure is perhaps more sensitive to experimenter demand effects than our other task-based measures but presents a consistent picture with them.

Application data from online search platform: Finally, we find that treated jobseekers alter their actual search direction on our partner's SAYouth.mobi platform. The platform aggregates job advertisements from many sources and offers data-free access, so jobseekers in our sample describe it as a key part of their job search strategy. Jobseekers in our sample actively used the platform in the 30 days prior to treatment, visiting it on 3.2 of those days and initiating 6.8 job applications (Table A20). We observe job advertisement titles and descriptions, which we use to classify job advertisements as numeracy or communication jobs, based on a list of common numeracy and communication jobs on the platform.[22] On average 23% of applications sent by jobseekers in the 30 days after the workshop are to jobs classified as either numeracy and communication.

Treated jobseekers are more likely to initiate applications to jobs that match their measured comparative advantage. We examine the difference between the fraction of applications that are

[22]We construct this list from (1) the job titles classified by the same 13 human resource and recruitment staff who classified advertisements for the job choice task and (2) job descriptions that explicitly require communication or numeracy skills. The former approach mimics the skill classification surveys used by O*NET, but is tailored to the specific skills and jobs relevant for this research. The latter approach follows existing work studying skill classification on platforms (e.g. Deming and Kahn 2018; Hershbein and Kahn 2018). Of the 69,000 vacancies available on SAYouth.mobi in Johannesburg during our study period, we classify 14% as communication jobs and 13% as numeracy jobs. This suggests that demand for skills is roughly equally distributed across the two skills. We observe if a jobseeker initiated an application by clicking on the "Apply" button. On average, jobseekers become more active after the workshop and initiate applications to 15 jobs on the platform in the following 30 days.

coded to be aligned with jobseekers' comparative advantage in skills and the fraction of applications coded as requiring the opposite skill. This measure thus captures both changes in the search direction towards the comparative advantage and changes away from jobs that require the skill without a comparative advantage. Treated jobseekers submit 6.3 percentage points more aligned applications than non-aligned applications ($p = 0.01$). Again, this effect is driven by jobseekers with initially misaligned comparative advantage beliefs (8.9 percentage points, $p = 0.072$) compared to jobseekers with initially aligned comparative advantage beliefs (4.1 percentage points), though the difference between the two is not significant. Because participants' information was obtained after the experiment from administrative data, we consider the platform data to be least susceptible to experimenter demand effects.

Search direction index: We combine these measures to an index to avoid multiple hypothesis testing and to increase power (Anderson, 2008). We see a strong, positive, and significant treatment effect on the index of our four measures of directed search of 0.27 standard deviations ($p = 0.013$). This effect is entirely driven by jobseekers with initially unaligned comparative advantage beliefs, for whom we observe a shift in search direction of 0.6 standard deviations ($p < 0.001$). The effect for jobseekers with initially aligned beliefs is close to zero (-0.05 standard deviations and not significant). Importantly, the difference in treatment effects between the two groups is significant at the one percent level (-0.65 standard deviations, $p = 0.003$).

The size of the search direction effect is similar to the correlation of aligned baseline beliefs with search direction in the control group (0.65 standard deviations, $p < 0.001$). The strong correlation in the control group further supports the notion that it is indeed the belief about the comparative advantage that drives measured search direction in this context. It is also consistent with the evidence from the field experiment discussed in Section 2.

3.5 Beliefs About Returns to Directed Search

The model predicts that changes in search direction will be driven by changes in beliefs about the relative returns to searching for jobs that require different skills, i.e., beliefs about the returns to skill-directed job search. To evaluate this prediction, we measure jobseekers' beliefs about the outcomes from applying to communication- and numeracy-heavy jobs, both in the job choice task during the workshop and in their planned search after the workshop. We measure the believed return to skill-directed search as the difference in expected outcomes between CA-aligned and CA-nonaligned jobs, e.g., for a jobseeker with numeracy CA, the expected wage for numeracy-heavy jobs minus the expected wage for communication-heavy jobs, averaging over all job pairs.

Treatment increases the believed returns to most measures of skill-directed job search. Furthermore, jobseekers' beliefs about job-specific expected wages predict which jobs they choose in the job choice task. These results are consistent with the model, although not all are precisely estimated. See Appendix B for details of the the measurement and results.

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

Having demonstrated that information about skill assessment results can shift search direction, we now test if this information can also improve individual labor market outcomes. We use a field experiment with 4,389 participants and a follow-up period of 3.5 months on average, giving us more power and time to study effects on labor market outcomes. The same experiment is used by Carranza et al. (2022) in related research.[23]

4.1 Sample

The big experiment took place in the same labor market in 2016/7, five years prior to the tight experiment. Recruitment, described in Section 2.1, is very similar to the tight experiment. The sample's characteristics are also fairly similar. The main difference is that jobseekers in this experiment are slightly younger than in the tight experiment, with an average age of 24 instead of 26 (Table 1, panels A and B). But they are similarly likely to face limited information about their skills in the labor market. Less than half have information on their skills from post-secondary education: 39% had attended post-secondary schooling (16pp higher than the tight experiment). Their formal work experience is relatively limited: 9% had ever held a permanent or long-term job (lower than the 23% in the tight experiment, perhaps linked to their younger age and greater level of education). Jobseekers were similarly actively engaged in the labor market: at baseline 37% (compared to 33% in the tight experiment) had done some work or income-generating activity in the past seven days and 97% (compared to 96%) were actively searching for work. The average jobseeker submitted 9 (compared to 10) job applications and spent 17 (compared to 14) hours searching for work in the last seven days. At baseline, they spend on average 31.63 USD per week on search expenses (compared to 22.72 USD in the tight experiment, but not adjusted for inflation). These substantial search efforts are, on average, associated with only 0.80 job offers in the last 30 days.

The baseline is a self-administered questionnaire that candidates complete on desktop computers at Harambee under supervision. This is administered after candidates have done skills assessments but before they receive information about their results. We collect follow-up data using 25-minute phone surveys on average 3.5 months after treatment, reaching 96% of the respondents.[24] The response rate is balanced across treatment groups (Table A23). We collect data on beliefs, search, and labor market outcomes of participants both at the baseline and at the follow-up.

[23] That paper studies the effect of providing information about jobseekers' skills to both firms and jobseekers, relative to providing it to only jobseekers or to neither, using multiple experiments. It focuses more on firm-side responses but does report the effects of giving jobseekers' their assessment results on CA-aligned search and some labor market outcomes. Our work adds comparative advantage beliefs, both as an outcome and driver of heterogeneous treatment effects, and effects on additional labor market outcomes.

[24] Garlick et al. (2020) provides an experimental validation of phone-based labor market surveys in this setting. Respondents receive mobile phone airtime payments for answering the phone surveys.

4.2 Experimental Design

Like the tight experiment, the big experiment also randomly varies whether jobseekers receive a report about their comparative advantage in skills. Treatment is administered at the day level. We assign 2,114 jobseekers assessed over 27 days to the treatment group and 2,274 jobseekers assessed over a different 27 days to the control group. Treatment groups are balanced on covariates both for the full and the sample recontacted for the endline survey (Table A24).

The design of the field experiment is very similar to that of the tight experiment. Jobseekers arrived at the Harambee offices, also in downtown Johannesburg, took skills assessments, took a self-administered baseline survey, and received some job search assistance in a workshop. As in the tight experiment, jobseekers in the treatment group receive an unbranded report that provides information about their performance on the skills tests they took during their visits at Harambee. The control group does not receive any information about their skills. They receive the business-as-usual experience at Harambee and were not told that jobseekers assessed on other days received assessment results. Therefore, it is unlikely that jobseekers assigned to the control group were discouraged by the lack of assessment results. Both groups knew they were participating in a research study but did not know the goals of the study or their own treatment status, so social desirability bias in their survey answers is unlikely.

The field experiment differs from the tight experiment in four ways. First, we assess jobseekers on six skills: communication, numeracy, concept formation, focus, grit, and planning, as described in Appendix D. And the reports given to treated jobseekers show results on all six assessments (Figure A3), rather than the two assessments in the tight experiment. Using more skills has the advantage of allowing jobseekers to direct their search across a broader range of job types, taking advantage of the substantial heterogeneity in firms' skill demands in this context, documented by Carranza et al. (2022).

The six assessments horizontally differentiate jobseekers based on their relative skills rather than only ranking or vertically differentiating them in a single dimension of skills, allowing skill-directed search in many different directions. Assessment results are weakly correlated across skills within candidates, with 12 of the 15 pairwise correlations below 0.2 (Table A43). As a result, most jobseekers learn they have substantial variation across skills: 85% have at least one top tercile but only 1.7% have six top terciles, 73% have at least one bottom tercile but only 0.5% have six bottom terciles, and 58% have both top and bottom terciles.

Second, we report assessment results in terciles and not quintiles to allow quick administration with larger groups of jobseekers (on average 80 per day). The coarseness of terciles relative to quintiles and using six rather than two skills means that only 23% of the jobseekers have a unique comparative advantage (CA). Hence restricting the sample to jobseekers with unique CA, as we did in the tight experiment, is not sensible. The tight and big experiments thus deliberately capture contrasting approaches: the tight experiment uses fewer skills to allow cleaner definitions of CA and clearer tests for skill-directed search, while the big experiment uses more skills to allow more skill-differentiated job search and hence potentially larger effects on labor market outcomes.

Third, the briefing after candidates receive the report differs slightly. Instead of treated candidate watching a video, in the big experiment they receive their reports during a group briefing with a psychologist. However, the content of briefings is highly consistent: we use the script from the psychologist's briefing to develop the video script for the tight experiment. The psychologist briefing explains what each assessment measures and how to interpret the results on the report. As in the video, the psychologist suggested potential uses of the information, including that the jobseekers should think about what this report suggests about the kinds of jobs they should to apply to.

Finally, the structure of the day differs slightly. Jobseekers receive a CV template, interview tips, and job search tips 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:

$$Y_{id} = T_d \cdot \beta + \mathbf{X}_{id} \cdot \Gamma + \varepsilon_i \tag{6}$$

where Y_{id} is the outcome of interest for jobseeker i assessed on date d and $\mathbf{X}_{id}$ is a vector of prespecified control variables.[25] The object of interest is the average treatment effect β. We cluster standard errors at the level of treatment (assessment days). As in the tight experiment, we combine families of outcomes in Anderson (2008) indices to limit the number of tested hypotheses. We show in Table A11 that our main hypothesis test results are robust to using a wild cluster bootstrap and to adjusting for multiple hypothesis testing using sharpened q-values, following Benjamini et al. (2006), while the main treatment effect estimates are robust to reweighting the field experiment sample to have the same distribution of baseline covariates as the tight experiment sample.

We do not emphasize heterogeneous treatment effects by baseline aligned CA beliefs as much as we did in the tight experiment. In the big experiment, with six skills, there are many more ways for assessed and believed CA to differ. Furthermore, we only observe beliefs about three of the six skills in the big experiment. Hence, not having baseline aligned CA beliefs is a noisier proxy for having beliefs that can be shifted by treatment in the big than tight experiment. While we emphasize heterogeneous treatment effects less in the big experiment, we still briefly discuss these results in the paper and show details in the appendices.

[25] $\mathbf{X}_{id}$ contains measured skills, self-reported skills, education, age, gender, employment, above median discount rate, above-median risk aversion, baseline values for the outcome where available, and fixed effects for the blocks of days within which treatment was randomized. This is identical to the list of controls in the tight experiment, except for a few baseline measures collected in only of the two experiments. The big experiment is preregistered at https://doi.org/10.1257/rct.1631-8.0. We describe the relationship between the preanalysis plan and our final analysis in Appendix F.

4.4 Better Information About Skills Aligns Beliefs with Comparative Advantage

Receiving information about skills substantially increases alignment between measured skills and skill beliefs at the time of the endline survey. This is in line with the effects on beliefs observed straight after treatment in the tight experiment. In the endline survey, we ask candidates in which tercile they believe they ranked for each of the communication, concept formation, and numeracy assessments.[26] We define jobseekers as having aligned CA beliefs if their believed top skill(s) matches their assessed top skill(s) across the three skills. If jobseekers have tied top skill terciles, we require them to also believe that they have tied top skill terciles. Only 19% of control group jobseekers have aligned CA beliefs on this measure. Treatment increases this by 13.9pp, a 72% increase ($p < 0.001$, column 1 of Table 7).[27] Treatment also increases fraction of skill beliefs that align with assessment results by 14.2pp, a 37% increase from the control mean of 38.8pp ($p < 0.001$, column 2). These results imply that the treatment worked as intended and aligned measured beliefs about skills with the assessment results.[28] As in the tight experiment, we observe no gender heterogeneity in baseline beliefs or belief updating, so we report all analysis pooling genders (Table A25).

4.5 Job Search with Better Aligned Beliefs about Skills

We collect a self-reported measure of whether jobseekers are directing their search according to skill CA. We ask jobseekers "Think about the types of jobs that you are applying for. Think about what skills are important for these jobs. Rank the following skills from most important to least important." The skill options are numeracy/maths skills, literacy/communications skills, problem-solving skills, and soft skills. We define a dummy for aligned search as equal to one if the jobseeker ranks (one of) the skills on which they perform best higher than the two other considered skills (ignoring soft skills). In other words, this variable is equal to one if the most important skill listed for jobs that they are applying for matches the skill in which they have an assessed CA.[29]

[26] The question wording is "Remember the assessments you took at Harambee. 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. Do you think you are in the top third, middle third or bottom third of people on the X test?" for each skill X.

[27] The treatment effects on aligned CA beliefs are almost identical across the tight and big experiments. However, the control group mean is substantially higher in the tight than big experiment because the big experiment uses three skills, creating more ways for believed and assessed CA to differ. In both experiments, the control group's alignment rate is only slightly higher than random guessing, accounting for the possibility of ties: 19 versus 15% in the big experiment and 48 versus 40% in the tight experiment.

[28] These questions capture jobseekers' *beliefs about their assessment results*, while the main questions in the tight experiment capture jobseekers' *beliefs about their general skills*. In the field experiment, we also asked beliefs about general skills for a random subsample of the first half of the jobseekers' surveyed, to allow a robustness check. Using these general skills, treatment increases aligned CA beliefs by 7.6pp ($p = 0.003$) and 10.8pp ($p < 0.001$) - smaller than the effects on beliefs about assessment results but qualitatively similar.

[29] For individuals without a clear CA (i.e., individuals with a tie for highest skill) our main aligned search measure is set to zero. However, all our results are robust to recoding this measure to one when any of the highest skills is listed first. We deliberately construct this measure of skill-aligned job search from separate questions about jobseekers' skill

Table 7: Treatment Effects on Skill Beliefs and Search Direction - Big Experiment

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

Notes: **Table 7 shows that providing jobseekers with information about their comparative advantage in skills aligns skill beliefs with assessed skills and shifts self-reported search direction in the big experiment.** Column 1 shows the impact on aligned comparative advantage beliefs. Column 2 shows the impact on the fraction of aligned skill tercile beliefs across three domains (numeracy, communication, and concept formation). Column 3 shows the impact on a dummy indicating search direction aligned with comparative advantage in skills (self-reported). Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). Standard errors clustered at the treatment-day level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Treatment increases the share of jobseekers searching mainly for jobs that match the skill in which they score highest from 17% to 22% (column 3, $p < 0.001$), three months after treatment. This self-reported search direction result is qualitatively consistent with the four measures of directed search in the tight experiment, collected within the first month after treatment. But the effect in the big experiment converts to a 0.14 standard deviation increase in directed search, compared to the 0.27 standard deviation increase in the directed search index in the tight experiment (Table 6, column 1). The latter effect is likely to be larger because the index averages over four measures, producing a smaller standard deviation and hence larger standardized effect size.

Taken together, the results from the two experiments provide strong support for our key argument, that information about skills improves the alignment of jobseekers' job-search strategy with their measured skill CA. This is in line with the predictions of the model that a shift in CA beliefs should change the search direction of jobseekers toward jobs that require the skills in which they perform well.

4.6 Improved Labor Market Outcomes

The model predicts that the change in search direction should then affect jobseekers' labor market outcomes as they direct their search toward jobs with higher returns to search. Here we show

beliefs and the skill demand of jobs to which they apply, rather than asking them directly if they are searching for jobs that match their CA, because the latter approach might be more susceptible to experimenter demand effects.

evidence consistent with this prediction and then evaluate possible mechanisms that link skill-directed search to 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 month after the workshop, second month after the workshop, and week before the endline survey are positive but mostly small and mostly not statistically significant (Table 8, columns 2-5). Our prespecified Anderson (2008) index of work quantity increases by 0.05 standard deviations with $p = 0.13$ (column 1).

Treatment has substantial effects on employment quality: Treatment increases earnings in the seven days before the endline by 6.52 USD, 26% of the control group mean (column 7, $p = 0.019$). 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. Interpreting treatment effects on such measures can be complex (Mullahy and Norton, 2022). Hence we show that treatment effects on earnings are robust across multiple different approaches. First, effects are weakly positive on the entire earnings distribution (Figure A4). Second, treatment raises earnings by 9.4 USD without winsorization, 0.1 standard deviations or 12 log points or 13 inverse hyperbolic sine points, demonstrating robustness across a range of transformations (Table A26, columns 1 and 3-5). Third, restricting the sample to the employed produces only slightly larger treatment effects (Table A26, columns 6-10). This third approach can create a sample selection problem when treatment affects the probability of employment and hence the probability of having non-missing earnings. But, in this case, treatment has little impact on earnings, so it is unsurprising that the results are robust to different ways of handling the earnings of the nonemployed. The earnings effects are at the upper end but still in line with effect sizes in related research.[30]

We show in Appendix E that the earnings effect implies that the average treated jobseeker earns enough extra in the time between treatment and endline to cover 1.8 times the average variable cost of the assessment operation. We discuss possible implications of this cost-benefit comparison in the conclusion.

Treatment increases hourly wages by 0.30 USD or 24 percent of the control mean of 1.27 USD (column 8 of Table 8, $p = 0.028$). Treatment also increases the probability of having a written contract, an important and widely desired marker of job formality in this context, by 1.7 percentage points from a 12% mean in the control group (column 9, $p = 0.095$). Our index of work quality increases by 0.085 standard deviations (column 6, $p = 0.019$).

Heterogeneous treatment effects: As our model suggests, these treatment effects are driven by jobseekers who get more new information from treatment about the comparative advantage. Table A29 shows that the treatment effect on every labor market outcome is larger for jobseekers with misaligned than aligned CA beliefs at baseline. Although none of the differences is statistically

[30]For example, interventions that combine learning about skills with the ability to signal skills increase earnings by 11-34% of the control group means (Abebe et al., 2021a; Bassi and Nansamba, 2022; Carranza et al., 2022). Bandiera et al. (2021) find an 11% change in earnings from a matching intervention that they attribute to changes in 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 wages by 11 percent.

Table 8: Treatment Effects on Labor Market Outcomes - Big Experiment

	(1)	(2)	(3)	(4)	(5)
Panel A: Work quantity	Index	Job offers (w)	Worked month 1	Worked month 2	Worked last 7 days
Treatment	0.045	0.015	0.023*	0.007	0.009
	(0.032)	(0.017)	(0.013)	(0.015)	(0.013)
Control mean	-0.000	0.182	0.465	0.437	0.309
Observations	4205	4140	4201	4204	4204
Panel B: Work quality	Index	Earnings (w)	Hourly wage (w)	Written contract	
Treatment	0.085**	6.517**	0.295**	0.017*	
	(0.035)	(2.712)	(0.130)	(0.010)	
Control mean	0.000	25.424	1.267	0.120	
Observations	4206	4196	4183	4184	

Notes: **Table 8 shows that the treatment improves work quality but not necessarily work quantity in the big experiment.** Panel A shows impacts on employment quantity. Panel A, column 1 shows impacts on an Anderson (2008) index of the four employment quantity measures. Panel A, column 2 shows the impact on the number of job offers in the last 30 days (winsorized). Panel A, column 3 shows impacts on a dummy indicating any work for pay in month 1 after treatment. Panel A, column 4 shows impacts on a dummy indicating any work for pay in month 2 after treatment. Panel A, column 5 shows the impact on a dummy indicating any work for pay in the last seven days. Panel B show impacts on employment quality. Panel B, column 1 shows impacts on an Anderson (2008) index of the three employment quality measures. Panel B, column 2 shows the impact on earnings in the last seven days (winsorized). Panel B, column 3 shows the impact on hourly wages in the last seven days (winsorized). Panel B, column 4 shows the impact on a dummy indicating a written contract. Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). All monetary figures are reported in 2021 USD PPP. All winsorized variables are winsorized at the 99^{th} percentile. Standard errors clustered at the treatment-day level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

significant at conventional levels, some are large; for example, the effects on weekly earnings are 7.5 and 2.8 USD for jobseekers with misaligned and aligned beliefs respectively. Similarly, the treatment effect on skill-directed job search is driven by jobseekers with misaligned baseline CA beliefs, although the pattern of heterogeneity is less clear for updating the CA beliefs themselves (Table A28).

Linking skill-directed search to improved labor market outcomes: What specific economic mechanisms link skill-directed search to 'better' jobs – in terms of earnings, hourly wages, and formality – but not to a higher employment probability? Our data do not allow a conclusive

answer to this question. But we provide a framework for thinking about these mechanisms and empirical tests of parts of this framework.

We begin by noting three patterns that motivate our focus on search-based mechanisms. In each case, we estimate treatment effects on interactions between multiple labor market outcomes, e.g., the treatment effects on earnings interacted with indicators for different start dates of current jobs. First, the majority of the positive effects on earnings and written contracts are in jobs with start dates after treatment (Table A27, columns 1-2 and 6-7). This suggests that changes to search behavior are more important than wage bargaining with existing employers. Second, the entire earnings effect comes from wage employment, not self employment (columns 4-5). This suggests that changes to job application behavior and application outcomes are an important mechanism. Third, search and employment are uncorrelated in the control group and treatment has effects of less than 1pp on all four possible search $\times$ work states. This suggests that the distinction between on-the-job and off-the-job search actions is not an important consideration, although the outcomes of the on- and off-the-job search may obviously differ.

Given these patterns, we evaluate two mechanisms centered on two standard concepts in the search literature: the **wage offer distribution** and **job offer probability**. Returning to our model from Section 2.5, we assume that there are multiple types of jobs 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 allocation of search effort across different types of jobs generates a wage offer distribution $F(w)$ and job offer probability P that is a weighted average across the different job types. As in standard job models, any jobseekers who receives an offer accepts it if the wage is above their reservation wage. Treatment, by shifting search toward job types that use jobseekers' comparative advantage, can shift F to the right and/or raise P.

In the first potential mechanism, skill-directed job search shifts at least part of $F(w)$ to the right, increasing average earnings conditional on employment. If the entire wage offer distribution shifted to the right, then more jobseekers would receive offers above their reservation wage, leading to higher employment. To match our null effect on employment, this mechanism requires an additional feature: either an increase in the reservation wage, so some of the higher wage offers are rejected; or that the shift in F occurs only above reservation wages. We do not observe the terms of job offers, so we cannot directly test that part of the mechanism. But the treatment effect on the CDF of earnings suggests a broad rightward shift in F (Figure A4). Treatment increases reservation wages in both experiments but the effects are under 2% of the control group mean and not statistically significant. We conclude that the data don't offer strong evidence for or against this mechanism.

In the second potential mechanism, skill-directed search raises P, leading to more job offers. To match the positive effect on employment, some jobseekers must receive multiple offers, so the additional offers allow them to select higher-wage offers or bargain for higher wages using their other offer? and hence earn more. To match the null effect on employment, treatment must also increase reservation wages, so most of the additional offers are rejected. We see some evidence

against this explanation: treatment effects on earnings and hourly wages are substantially larger than the near-zero effects on reservation wages, the treatment effect on job offers is close to zero, and only 4% of jobseekers in both the control and treatment groups received more than one job offer in the month before endline. We conclude that the data provide more evidence against than for this mechanism.

If treatment does lead to higher wage offers, as in the first economic mechanism, why might it do so? The model suggests that jobseekers apply to jobs that better match their skills where they can be more productive. We can't directly test this in the big experiment, because we don't observe detailed data on occupation or job tasks. But we do find that treatment increases tenure by 0.44 months, a substantial increase given the control group mean of 1.9 months and the average 3.5 month gap between treatment and endline. This higher tenure provides suggestive but imprecise evidence for higher productivity due to better skill match.

Taken together, the treatment effects across the two experiments align with the model prediction that jobseekers' beliefs about their comparative advantages matter for search direction and, ultimately, labor market outcomes. The tight experiment provides stronger evidence for the link between comparative advantage beliefs and search direction, while the big experiment shows the labor market implications of this process.

5 Additional Mechanisms

Across the tight and big experiments, average and heterogeneous treatment effects are consistent with a skill-directed search interpretation: jobseekers receive information about their skills, shift comparative advantage beliefs toward assessed comparative advantage, and redirect search toward jobs that demand skills matching their comparative advantage, raising earnings and wages.

In this section, we consider five other mechanisms that might account for effects on earnings and wages. These are not in principle mutually exclusive with the directed search interpretations, as multiple mechanisms might link the treatments to labor market outcomes. In practice, however, we find little evidence for any of these five mechanisms.[31]

5.1 Search Effort and Beliefs About Skill Levels

In this subsection, we show that treatment updates jobseekers' beliefs about the level of their skills, in addition to the skill comparative advantage we focused on earlier, then generalize our conceptual framework to show that this update has a theoretically ambiguous effect on total search effort, and finally show that treatment has limited effects on search effort in both experiments. We conclude that search effort and beliefs about skill levels are unlikely to explain the labor market effects we find.

Treatment effects on beliefs about skill levels: Treatment shifts jobseekers' beliefs about their skill levels toward their measured skill levels. Treated jobseekers in the tight experiment reduce

[31] We prespecified testing for the "search effort" and "human capital investment" explanations in the tight experiment and for the "search effort," "self-esteem," "human capital investment," and "information transmission to firms" explanations in the big experiment.

their believed skill level by an average of 0.08 quintiles over the two skills ($p = 0.025$, Table A30, column 3). Treated jobseekers in the big experiment reduce their believed skill level by an average of 0.11 terciles over the three skills ($p < 0.001$, Table A30, column 7). In both samples, more jobseekers have baseline skill beliefs above their assessed skills than below. Hence more jobseekers receive downward than upward signals, explaining the negative average treatment effects on believed skill levels.

Treatment effects vary by baseline beliefs about skill levels. To show this, we construct a measure of baseline confidence levels: believed skill level minus assessed skill level, using quintiles and averaging over two skills in the tight experiment and using in terciles and averaging over three skills in the big experiment. The treatment effects on endline belief levels are substantially and statistically significantly lower for jobseekers with higher baseline confidence levels (Table A30, columns 4 and 8, row 2).

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 2.5 with the assumption that jobseekers choose the levels of search for both communication and numeracy jobs, E_C and E_N. This gives a utility function with three arguments: the expected outcome of search for communication jobs, $V_C(S_C, S_N, E_C)$, the expected outcome of search for numeracy 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: $A = Y - P_S \cdot E_C - P_S \cdot E_N$, where Y is the jobseeker's unearned income, and P_S 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

$$\max_{E_C, E_N} U\left(V_C(S_C, S_N, E_C), V_N(S_C, S_N, E_N), Y - P_S \cdot E_C - P_S \cdot E_N\right). \tag{7}$$

As in Section 2.5, we continue to 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 rise in the believed level of communication skill S_C raises the believed marginal product of search for communication jobs, $\frac{\partial V_C}{\partial E_C}$. This has two effects. First, a substitution effect, which causes the jobseeker to substitute toward search for communication jobs and away from both search for numeracy jobs and alternative activities. Second, an income effect: it raises the expected outcome from any given level of search effort, allowing the worker to allocate more money to alternative uses, and hence reducing search for both communication and numeracy jobs. The net effect is a decrease in search for numeracy jobs and an ambiguous effect on search for communication jobs, and hence an ambiguous effect on total search effort.

It may be helpful to note that this problem has the same structure as the standard static labor supply model. In that model, a rise in the hourly wage increases labor supply by raising the

opportunity cost of leisure (substitution effect) and decreases labor supply because the worker can afford more leisure (income effect), leading to a theoretically ambiguous total effect on labor supply. In our model, a rise in the marginal product of search due to a higher perceived skill level increases search by raising the opportunity cost of other spending (substitution effect) and decreases search because the worker can afford to spend more on other activities (income effect), leading to a theoretically ambiguous total effect on search effort. Abebe et al. (2022) also show that raising expected job search outcomes has an ambiguous effect on search effort using a frictional matching model.

This framework implies that treatment effects on search effort are theoretically ambiguous for the average jobseeker and for jobseekers who receive downward or upward signals from the experiment. We should not expect that treatment effects will have opposite signs for jobseekers who received downward and upward signals, as the income effect might exceed the substitution effect for one group and vice versa for the other group.

Treatment effects on search effort: We find little evidence of treatment effects on search effort in either experiment. In the tight experiment, average treatment effects are negative on five of six different measures of search effort but all effects are small – less than 10% of the control group mean – and none is statistically significant (Tables A31 and A32). Treatment lowers an index of these search measures by 0.08 standard deviations ($p = 0.47$). Treatment effects do not vary substantially by jobseekers' baseline confidence levels.[32]

In the big experiment, the average treatment effect on three search effort measures – applications submitted, hours and money spent searching – are positive but tiny ($< 3\%$ of the control group mean) and, again, none is statistically significant (Table A33). Treatment has a negligible effect on an index of these search measures: 0.003 standard deviations ($p = 0.92$). As in the tight experiment, treatment effects do not vary substantially by jobseekers' baseline confidence levels. The treatment effects on search effort have different signs in the two different experiments but all are close to zero and we cannot reject equality of effects on the two indices ($p = 0.47$). For interested readers, we provide a more detailed comparison of search effort effects across the two experiments in Appendix A.4.

We conclude that a causal chain based on skill belief levels and search effort is unlikely to explain the labor market treatment effects. Despite the positive effects on multiple job quality measures, treatment does not increase any search effort measure in either experiment, nor does it shift patterns of on-the-job search in the big experiment. This does not, of course, imply that search effort plays no role in determining labor market outcomes in this or other settings. It merely suggests that shifts in search effort are less likely to explain the treatment effects on labor market

[32]Our six search effort measures are planned applications from 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 after the workshop: days active, jobs viewed, and applications submitted. The planned applications, text messages, job applications are the same search measures 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 cover letter for a real job application at the end of the workshop, on a laptop we provided, rather than collecting their incentive and leaving early.

outcomes than shifts in search direction to align with skill comparative advantage.

5.2 Self-Esteem

Treatment might affect labor market outcomes if jobseekers' self-esteem responds to new skill information. Measures of self-esteem capture general beliefs about self-worth, rather than the narrower beliefs about skill we study in the rest of the paper, but might also affect search behavior.

Treatment has near-zero effects on self-esteem in the big experiment (Table A35). We find this result both in a text message survey 2-3 days after treatment (columns 1-2) and in the endline phone survey roughly 3.5 months after treatment (columns 3-4). The endline survey uses a five-item scale adapted from the Rosenberg (1965) scale and the text message survey uses a single item from this scale, all with five-point Likert response options.

5.3 Human Capital Investment

Treatment might affect labor market outcomes if jobseekers respond to new skill information by investing in new skills or shifting their investments between different types of skills, as documented in some education research (Stinebrickner and Stinebrickner, 2014b). We test this in both experiments. In the big experiment, treatment has near-zero effects on enrollment in formal and vocational education in the time between treatment and endline (columns 5 to 7 of Table A35). In the tight experiment, we measure treatment effects on jobseekers' willingness-to-pay (WTP) for numeracy and communication training materials. Treatment reduces 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. See Table A42 for results and Appendix C for details of the WTP measurement.

We conclude that the treatment effects on earnings in the big experiment are unlikely to be driven by changes in education investment. However, results in the tight experiment suggest that jobseekers might invest in skills under other conditions. For example, jobseekers might enroll in formal education over longer time periods or in environments with cheaper education.

5.4 Skill Information Transmission to Firms

Treatment might affect labor market outcomes if jobseekers share assessment results with firms during job applications, leading to firm-side learning about jobseekers' skills. Some jobseekers in the big experiment do share assessment results with firms. 29% of treated jobseekers surveyed at endline report including a copy of their assessment result with at least one job application, although only 0.8% of treated jobseekers actually included their assessment result when applying to a vacancy we created for a companion study.

However, sharing assessment results with firms is unlikely to explain the treatment effects on labor market outcome, for two reasons. First, the reports showing assessment results are deliberately designed not to be credible to firms. They do not show the jobseeker's name or national

identity number, so the firm cannot verify that the report is linked to that job applicant. They include no information about Harambee and no information designed to "sell" the assessment results, such as the source of the assessments or information about the value of skills we discuss in Section 2.3. We interviewed 15 hiring managers while piloting the study and none of them reported that they would view these reports as credible. Second, we estimate treatment effects on labor market outcomes separately for jobseekers who do and do not say that they use the reports in job applications. The earnings effects are driven by those who do not use the report, suggesting that information transmission to firms does not explain the average treatment effects on job quality (Table A36).[33] However, we interpret this result cautiously because we are estimating heterogeneous treatment effects over a post-treatment outcome.

5.5 Feedback Loop from Labor Market Experiences to Skill Beliefs

Treatment might affect labor market outcomes through a feedback loop between beliefs and labor market experience. Specifically, treatment might shift beliefs about skill levels or comparative advantage, leading to shifts in employment that provide information that in turn shifts skill beliefs.

We find little evidence consistent with feedback loop. First, treatment has a minimal effect on employment, so this feedback loop would have to operate through the type, not quantity, of employment. Second, treatment in the tight experiment has large effects on skill beliefs before there's time for any learning from labor market experience. Third, employment in the big experiment control group is uncorrelated with changes in beliefs, as we discuss in Section 2.4.

Even if this mechanism were large, it would not undermine the skill-directed search explanation. It would simply suggest that the magnitude of the treatment effects on labor market outcomes reflects both a direct effect and an indirect effect through this feedback loop.

6 Conclusion

This paper shows that skill-directed job search can improve jobseekers' labor market outcomes but that it can be constrained by jobseekers' limited information about their own comparative advantage in skills. Giving jobseekers additional information about their 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.

We provide evidence of a relatively understudied type of job search friction: misdirected search due to limited information about the comparative advantage in skills. This is important for interpreting the established research on firms' limited information about jobseekers' skills and jobseekers' limited information about their labor market prospects. For example, misdirected job search can reduce the effectiveness of firms' investment in screening prospective jobseekers, because it means firms don't see the ideal pool of applicants. Similarly, jobseekers' stated beliefs about their labor market prospects implicitly condition on beliefs about their skills and their search

[33] See Carranza et al. (2022) for the labor market effects of providing a different type of skill assessment result that is deliberately designed to be credible to firms.

strategies given these skills, so directly capturing these can better understand how beliefs and search behavior co-evolve. Our findings can inform models of job search and matching with multidimensional skills by providing direct evidence supporting models in which jobseekers' limited information about their relative skills distorts how they direct search.

Our approach highlights the benefits of combining multiple types of experimental designs with multiple types of measurement. This allows us to trace the entire causal chain from shifts in beliefs to changes in search direction and ultimately improved labor market outcomes. Showing consistent results from two similar but non-identical experiments provides a partial test of the generalizability of our findings.

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 the context of research like this? And by whom? The average treated participant's earnings gain likely exceeds the average variable cost of running the assessment system, suggesting the possibility of profitable market provision. Assessment allows substantial economies of scale, particularly if implemented by online job search and matching platforms, rather than the in-person assessment service we study. Many platforms, including our partner's 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 in at least the general skills we study in this project, 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, which sometimes include skill assessments, could fill this gap. Researchers have not yet studied the role of this specific component of job search counseling. Future work might examine the economics information provision by both public- and private-sector actors about comparative advantage across different types of skills.

References

ABEBE, G., A. S. CARIA, M. FAFCHAMPS, P. FALCO, S. FRANKLIN, S. QUINN, S. CARIA, M. FAFCHAMPS, P. FALCO, S. FRANKLIN, AND S. QUINN (2021a): "Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City." *Review of Economic Studies*, 88, 1279–1310.

ABEBE, G., A. S. CARIA, AND E. ORTIZ-OSPINA (2021b): "The Selection of Talent: Experimental and Structural Evidence from Ethiopia," *American Economic Review*, 111, 1757–1806.

ABEBE, G., S. CARIA, M. FAFCHAMPS, P. FALCO, S. FRANKLIN, S. QUINN, AND F. SHILPI (2022): "Matching Frictions and Distorted Beliefs: Evidence from a Job Fair Experiment," *Working Paper*.

ABEL, M., R. BURGER, AND P. PIRAINO (2020): "The Value of Reference Letters: Experimental Evidence from South Africa," *American Economic Journal: Applied Economics*, 12, 40–71.

ACEMOGLU, D. AND D. AUTOR (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by D. Card and O. Ashenfelter, Elsevier, 1043–1171.

ALFONSI, L., O. BANDIERA, V. BASSI, R. BURGESS, I. RASUL, M. SULAIMAN, AND A. VITALI (2017): "Tackling Youth Unemployment: Evidence from a Labor Market Experiment in Uganda," Manuscript, University College London.

ALFONSI, L., M. NAMUBIRU, AND S. SPAZIANI (2022): "Meet Your Future: Experimental Evidence on the Labor Market Effects of Mentors," *Working Paper*.

ALTMANN, S., A. FALK, S. JÄGER, AND 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. GLENNY, R. MAHLSTEDT, AND A. SEBALD (2022): "The Direct and Indirect Effects of Online Job Search Advice," *IZA Discussion Paper 15830*.

ALTONJI, J. AND C. PIERRET (2001): "Employer Learning and Statistical Discrimination," *Quarterly Journal of Economics*, 116, 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, 1481–1495.

ARCIDIACONO, P., E. AUCEJO, A. MAUREL, AND T. RANSOM (2016): "College Attrition and the Dynamics of Information Revelation," *NBER Working Paper 22325*.

ARCIDIACONO, P., P. BAYER, AND A. HIZMO (2010): "Beyond Signaling and Human Capital: Education and the Revelation of Ability," *American Economic Journal: Applied Economics*, 2, 76–104.

BALEY, I., A. FIGUEIREDO, AND R. ULBRICHT (2022): "Mismatch Cycles," *Journal of Political Economy*, 130.

BANDIERA, O., V. BASSI, R. BURGESS, I. RASUL, M. SULAIMAN, AND A. VITALI (2021): "The Search for Good Jobs: Evidence from a Six-Year Field Experiment in Uganda," *SSRN Working Paper 3910330*.

BANDIERA, O., N. PAREKH, B. PETRONGOLO, AND M. RAO (2022): "Men Are from Mars, and Women Too: A Bayesian Meta-analysis of Overconfidence Experiments," *Economica*, 89, 38–70.

BANERJEE, A. V. AND S. SEQUEIRA (2020): "Spatial Mismatches and Imperfect Information in the Job Search," *CEPR Discussion Paper 14414*.

BASSI, V. AND A. NANSAMBA (2022): "Screening and Signalling Non-Cognitive Skills: Experimental Evidence from Uganda," *The Economic Journal*, 132, 471–511.

BAZZI, S., L. CAMERON, S. G. SCHANER, AND F. WITOELAR (2021): "Information, Intermediaries, and International Migration," *NBER Working Paper 29588*.

BELOT, M., P. KIRCHER, AND P. MUELLER (2022a): "Do the Long-term Unemployed Benefit from Automated Occupational Advice during Online Job Search?" *IZA Discussion Paper 15452*.

BELOT, M., P. KIRCHER, AND P. MULLER (2018): "Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice," *Review of Economic Studies*, 86, 1411–1447.

——— (2022b): "How Wage Announcements Affect Job Search – A Field Experiment," *American Economic Journal: Macroeconomics*, 14, 1–67.

BENJAMINI, Y., A. KRIEGER, AND D. YEKUTIELI (2006): "Adaptive Linear Step-Up Procedures That Control the False Discovery Rate," *Biometrika*, 93, 491–507.

CARRANZA, E., R. GARLICK, K. ORKIN, AND N. RANKIN (2022): "Job Search and Hiring with Limited Information about Workseekers' Skills," *American Economic Review*, 112, 3547–83.

CHEN, J. AND J. ROTH (2022): "Log-like? Identified ATEs Defined with Zero-Valued Outcomes Are (Arbitrarily) Scale-Dependent," *Working Paper*.

CONLON, J. J., L. PILOSSOPH, M. WISWALL, AND B. ZAFAR (2018): "Labor Market Search With

Imperfect Information and Learning," *NBER Working Paper 24988*.

CORTÉS, P., J. PAN, L. PILOSSOPH, AND B. ZAFAR (2021): "Gender Differences in Job Search and the Earnings Gap: Evidence from Business Majors," *NBER Working Paper 28820*.

DELAVANDE, A. (2022): "Expectations in Development Economics," in *Handbook of Economic Expectations*, ed. by W. Van der Klaauw, G. Topa, and R. Bachman, Elsevier, 261–291.

DEMING, D. AND L. KAHN (2018): "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals," *Journal of Labor Economics*, 36, S337–S369.

DIAMOND, A. (2013): "Executive Functions," *Annual Review of Psychology*, 64, 135–168.

DUCKWORTH, A. (2016): *Grit: The Power of Passion and Perseverance*, New York, NY, US: Scribner/Simon & Schuster.

FALK, A., D. HUFFMAN, AND U. SUNDE (2006): "Self-Confidence and Search," *IZA Discussion Paper 2525*.

FIELD, E., R. GARLICK, N. SUBRAMANIAN, AND K. VYBORNY (2023): "Why Don't Jobseekers Search More? Barriers and Returns to Search on a Job Matching Platform," *Working Paper*.

GARLICK, R., K. ORKIN, AND S. QUINN (2020): "Call Me Maybe: Experimental Evidence on Frequency and Medium Effects in Microenterprise Surveys," *The World Bank Economic Review*, 34, 418–443.

GIBBONS, R., L. KATZ, T. LEMIEUX, AND D. PARENT (2005): "Comparative Advantage, Learning, and Sectoral Wage Determination," *Journal of Labor Economics*, 23, 681–724.

GNEEZY, U., A. RUSTICHINI, AND A. VOSTROKNUTOV (2010): "Experience and Insight in The Race Game," *Journal of Economic Behavior and Organization*, 75, 144–155.

GUVENEN, F., B. KURUSCU, S. TANAKA, AND D. WICZER (2020): "Multidimensional Skill Mismatch," *American Economic Journal: Macroeconomics*, 12, 210–244.

HAALAND, I., C. ROTH, AND J. WOHLFART (2023): "Designing Information Provision Experiments," *Journal of Economic Literature*, 61, 3–40.

HARRISON, G. AND J. LIST (2004): "Field Experiments," *Journal of Economic Literature*, 42, 1009–1055.

HE, H., D. NEUMARK, AND Q. WENG (2021): ""I Still Haven't Found What I'm Looking For": Evidence of Directed Search from a Field Experiment," *NBER Working Paper 28660*.

HEATH, R., L. BOUDREAU, AND T. MCCORMICK (2022): "Migrants, Experience, and Working Conditions in Bangladeshi Garment Factories," *SSRN Working Paper 4067206*.

HERSHBEIN, B. AND L. KAHN (2018): "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings," *American Economic Review*, 108, 1737–1772.

HOFFMAN, M. AND S. BURKS (2020): "Worker Overconfidence: Field Evidence and Implications for Employee Turnover and Firm Profits," *Quantitative Economics*, 11, 315–348.

HUFFMAN, D., C. RAYMOND, AND J. SHVETS (2022): "Persistent Overconfidence and Biased Memory: Evidence from Managers," *American Economic Review*, 112, 3141–75.

JONES, S. AND R. SANTOS (2022): "Can Information Correct Optimistic Wage Expectations? Evidence From Mozambican Job-seekers," *Journal of Development Economics*, 159, 102–987.

KAHAN, D. M. (2013): "Ideology, Motivated Reasoning, and Cognitive Reflection," *Judgment and Decision Making*, 8, 407–424.

KAHN, L. AND F. LANGE (2014): "Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures," *Review of Economic Studies*, 84, 1575–1613.

KELLEY, E. M., C. KSOLL, AND J. MAGRUDER (2020): "How Do Online Job Portals Affect Employment and Job Search? Evidence from India," *Working Paper*.

KERR, A. (2017): "Tax(i)ing The Poor? Commuting Costs in South African Cities," *South African Journal of Economics*, 85, 321–340.

KESSLER, J. B., C. LOW, AND C. D. SULLIVAN (2019): "Incentivized Resume Rating: Eliciting

Employer Preferences without Deception," *American Economic Review*, 109, 3713–44.
KRUGER, J. AND D. DUNNING (1999): "Unskilled and Unaware of It: How Difficulties in Recognizing One's own Incompetence Lead to Inflated Self-Assessments." *Journal of Personality and Social Psychology*, 77, 1121.
LAM, D., C. ARDINGTON, AND M. LEIBBRANDT (2011): "Schooling as a Lottery: Racial Differences in School Advancement in Urban South Africa," *Journal of Development Economics*, 95, 121–136.
LENTZ, R., J. MAIBOM, AND 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. AND F. POSTEL-VINAY (2020): "Multidimensional Skills, Sorting, and Human Capital Accumulation," *American Economic Review*, 110, 2328–76.
MACLEOD, W. B., E. RIEHL, J. SAAVEDRA, AND M. URQUIOLA (2017): "The Big Sort: College Reputation and Labor Market Outcomes," *American Economic Journal: Applied Economics*, 9, 223–261.
MAHONEY, N. (2022): "Principles for Combining Descriptive and Model-Based Analysis in Applied Microeconomics Research," *Journal of Economic Perspectives*, 36, 211–222.
MALMENDIER, U. AND G. TATE (2015): "Behavioral CEOs: On the Role of Managerial Overconfidence," *Journal of Economic Perspectives*, 29, 37–60.
MAS, A. AND A. PALLAIS (2017): "Valuing Alternative Work Arrangements," *American Economic Review*, 107, 3722–3759.
MUELLER, A. I. AND J. SPINNEWIJN (2022): "Expectations Data, Labor Market and Job Search," in *Handbook of Economic Expectations*, ed. by R. Bachmann, G. Topa, and W. van der Klaauw, Netherlands: Elsevier Science, chap. 22, 677–710.
MUELLER, A. I., J. SPINNEWIJN, AND G. TOPA (2021): "Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence and Bias," *Econometrica*, 111, 324–363.
MULLAHY, J. AND 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*.
PALLAIS, A. (2014): "Inefficient Hiring in Entry-Level Labor Markets," *American Economic Review*, 104, 3565–3599.
PILLAY, A. L. (2020): "Prioritising Career Guidance and Development Services in Post-Apartheid South Africa," *African Journal of Career Development*, 2, 1–5.
POSNER, M. AND G. DIGIROLAMO (1998): "Executive Attention: Conflict, Target Detection, and Cognitive Control," in *The Attentive Brain*, ed. by R. Parasuraman, MIT Press, 401–423.
PRITCHETT, L. (2013): *The Rebirth of Education: Schooling Ain't Learning*, Washington, DC: Center for Global Development.
RAVEN, J. AND J. RAVEN (2003): "Raven Progressive Matrices," in *Handbook of Nonverbal Assessment*, ed. by R. McCallum, Boston: Springer, 223–237.
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, 135–146.
SANDERS, C. AND C. TABER (2012): "Life-Cycle Wage Growth and Heterogeneous Human Capital," *Annual Review of Economics*, 4, 399–425.
SANTOS-PINTO, L. AND 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, AND 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, 9–18.
SOCKIN, J. AND A. SOJOURNER (2020): "What's the Inside Scoop? Challenges in the Supply and Demand for Information about Job Attributes," *Working Paper*.
SPINNEWIJN, J. (2015): "Unemployed but Optimistic: Optimal Insurance Design With Biased Beliefs," *Journal of the European Economic Association*, 13, 130–167.
STATISTICS SOUTH AFRICA (2022): *Quarterly Labor Force Survey Quarter 3 2022*, Pretoria: Statistics South Africa.
STINEBRICKNER, R. AND T. STINEBRICKNER (2014a): "Academic Performance and College Dropout: Using Longitudinal Expectations Data to Estimate a Learning Model," *Journal of Labor Economics*, 32, 601–644.
——— (2014b): "A Major in Science? Initial Beliefs and Final Outcomes for College Major and Dropout," *Review of Economic Studies*, 81, 426–472.
STROOP, J. R. (1935): "Studies of Interference in Serial Verbal Reactions," *Journal of Experimental Psychology*, 18, 643–662.
SUBRAMANIAN, N. (2022): "Workplace Attributes and Women's Labor Supply Decisions: Evidence from a Randomized Experiment," *Working Paper*.
TAYLOR, S., S. VAN DER BERG, V. REDDY, AND D. J. VAN RENSBURG (2011): "How Well Do South African Schools Convert Grade 8 Achievement into Matric Outcomes?" *Stellenbosch University, Department of Economics Working Papers 13/2011*.
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.
WISWALL, M. AND B. ZAFAR (2015): "Determinants of College Major Choice: Identification Using an Information Experiment," *The Review of Economic Studies*, 82, 791–824.
WU, L. (2021): "Partially Directed Search in the Labor Market," *Working Paper*.
ZIMMERMANN, F. (2020): "The Dynamics of Motivated Beliefs," *American Economic Review*, 110, 337–363.

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

This contains eight appendices. Appendix A contains appendix tables and figures. Appendix B provides an additional analysis of the role of labor market beliefs in shaping search direction. Appendix C contains details and results of the willingness-to-pay elicitation in the tight experiment. Appendix D contains details about skill measurement. Appendix E details the cost-benefit calculation. Appendix F describes how our analysis relates to the pre-analysis plans. Appendix G describes the experimental protocols in detail. Appendix H describes how we measured firm demand for skills.

A Appendix Tables and Figures

The appendix tables are divided into four subsections. Appendix A.1 contains additional motivational exhibits, supporting section 2 of the paper. Appendices A.2 and A.3 contain appendix tables and figures that complement the tight and big experiment respectively, supporting sections 3 and 4. Appendix A.4 contains exhibits supporting the additional mechanism analyses from Section 5.

A.1 Motivational Appendix

This section contains tables that motivate the experimental design and analysis. Table A1 displays correlations between self-reported grades in English and math in the secondary school leaving exam (matric) and assessment results and beliefs. Table A2 shows the joint distribution of skill quintiles. Table A3 displays gender differences in baseline skill beliefs in the tight experiment. Table A4 shows unadjusted and adjusted gender differences in beliefs in the big experiment. Table A5 shows the development of skill beliefs over time in the control group of the big experiment. Table A6 shows the correlation between jobseekers' beliefs about skills and beliefs about the returns to search using the control group in the big experiment. Table A7 shows the correlation between jobseekers' beliefs about comparative advantage and search direction in the control group of the big experiment.

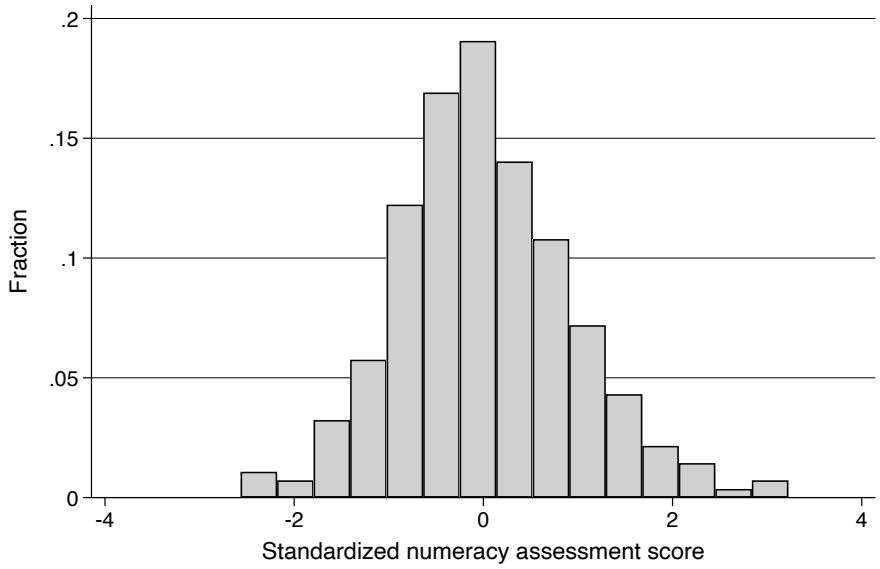
Table A1: Association Between High School Graduation Exam Results and Assessed and Beliefs About Skills - Tight Experiment

	Assessed skills			Beliefs about skills			
	Skill quintile		Comp. adv.	Skill quintile		Comp. adv.	
	(1) Num.	(2) Com.	(3) Num.	(4) Num.	(5) Com.	(6) Num.	(7) Com.
Matric: Math score	1.451***	0.594		1.782***	-0.360		
	(0.506)	(0.609)		(0.399)	(0.261)		
Matric: English score	0.219	1.219**		-0.355	1.016***		
	(0.400)	(0.477)		(0.335)	(0.223)		
Matric: Δ Math score - English score			0.192			0.308***	-0.487***
			(0.146)			(0.108)	(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

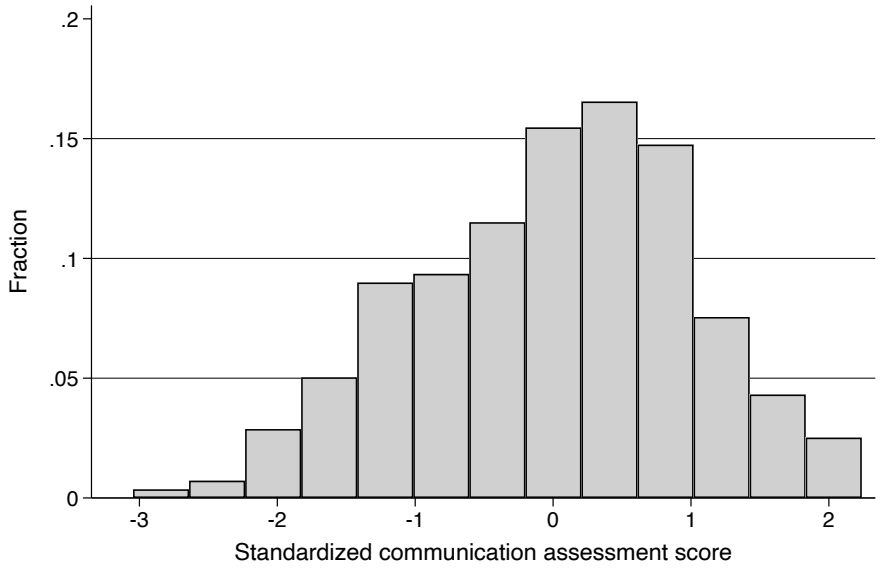
Notes: **Table A1 shows that self-reported grades in English and math in the secondary school leaving exam (matric) correlate positively with jobseekers' assessed skills and their baseline beliefs about skills.** It displays correlations between comparative advantage assessed based on matric scores and comparative advantage assessed based on assessment results and beliefs in the tight experiment. The sample includes all individuals with a clear comparative advantage who reported matric scores in English and math. Columns 1 to 3 correlate matric grades with assessment results. Columns 4 to 7 correlate matric grades with baseline beliefs about skills. Columns 3, 6, and 7 correlate the difference in math and English matric grades with comparative advantage measures. Columns 1, 2, 4, and 5 correlate math and English matric scores with assessed skill quintiles (columns 1 and 2) and beliefs about skill quintiles (columns 4 and 5). Matric grades are rescaled to range from 0 to 1. No further control variables are included. Robust standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure A1: Distribution of Measured Skills - Tight Experiment

Panel A: Numeracy



Panel B: Communication



Notes: **Figure A1 shows the distribution of standardized assessment scores in the tight experiment.** Panel A displays numeracy scores. Panel B displays communication scores.

Table A2: Joint Distribution of Skill Quintiles - Tight Experiment

Numeracy quintile	Communication quintile				
	Fifth	Fourth	Third	Second	First
Fifth	0.00	7.19%	6.47%	4.68%	4.68%
Fourth	11.51%	0.00	7.19%	7.91%	9.35%
Third	3.96%	2.52%	0.00	4.68%	2.88%
Second	2.52%	4.32%	3.96%	0.00	7.19%
First	0.72%	2.16%	0.72%	5.40%	0.00

Notes: **Table A2 shows that there is a lot of variation in quintiles across numeracy and communication skills in the tight experiment.** Sample is restricted to jobseekers with a unique comparative advantage in skills.

Table A3: Gender Differences in Skill Beliefs - Tight Experiment

	Female	Male	Δ	p($\Delta = 0$)	$\Delta(adjusted)$	p($\Delta(adjusted) = 0$)	N
Aligned comp. adv belief	0.50	0.46	-0.04	0.58	-0.01	0.83	278
Aligned comp. adv belief (assessment result)	0.56	0.48	-0.08	0.22	-0.03	0.46	278
Av. abs(beliefs - assessments) (general skill)	2.36	2.38	0.01	0.88	0.04	0.63	278
Av. abs(beliefs - assessments) (assessment result)	2.12	1.89	-0.23	0.03	-0.16	0.13	278
Fraction overcofident beliefs	0.62	0.59	-0.03	0.52	0.04	0.11	278
Fraction underercofident beliefs	0.15	0.20	0.05	0.15	0.01	0.79	278

Notes: **Table A3 shows that gender differences in beliefs in the tight experiment are small.** Adjusted gender differences control for age, education level, skill quintiles, and randomization block fixed effects.

Table A4: Gender Differences in Skill Beliefs - Big Experiment

	Female	Male	Δ	p($\Delta = 0$)	$\Delta(adjusted)$	p($\Delta(adjusted) = 0$)	N
Aligned comp. adv belief	0.19	0.23	0.04	0.01	0.03	0.03	4312
Fraction aligned beliefs	0.35	0.42	0.07	0.00	0.01	0.24	4378
Fraction overconfident beliefs	0.53	0.45	-0.09	0.00	0.00	0.75	4378
Fraction underconfident beliefs	0.11	0.13	0.01	0.06	-0.01	0.04	4378

Notes: **Table A4 shows that gender differences in beliefs in the big experiment are small**. Adjusted differences control for pre-specified covariates (measured skills, self-reported skills, education, age, employment, discount rate, and risk aversion).

Table A5: Development of Skill Beliefs Over Time - Big Experiment Control Group

	Aligned comparative advantage belief			% aligned skill belief		
	(1)	(2)	(3)	(4)	(5)	(6)
Endline	-0.004	-0.001	-0.015	0.009	0.009	0.004
	(0.011)	(0.016)	(0.013)	(0.008)	(0.011)	(0.010)
Endline × Above median search effort		-0.008			0.000	
		(0.025)			(0.014)	
Above median search effort		0.035*			0.003	
		(0.019)			(0.012)	
Endline × Worked last 7 days			0.034			0.015
			(0.022)			(0.014)
Worked last 7 days			-0.005			0.013
			(0.014)			(0.016)
Constant	0.200***	0.183***	0.202***	0.379***	0.377***	0.375***
	(0.009)	(0.014)	(0.010)	(0.007)	(0.007)	(0.009)
Observations	4405	4315	4315	4456	4365	4365

Notes: **Table A5 shows that the misalignment of skill beliefs persists over time in the control group of the big experiment.** Estimation is at the survey round times jobseeker level. Columns 1 to 3 show results for aligned beliefs about comparative advantage. Columns 4 to 6 show results for the fraction of aligned skill beliefs. Columns 2 and 5 show heterogeneity by whether individuals exerted above median search effort as measured on the search effort index described in Table A33. Columns 3 and 6 show heterogeneity by whether an individual worked in the last seven days. Heteroskedasticity robust standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A6: 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***	-0.129***	38.318***	33.755***
	(0.047)	(0.046)	(9.518)	(9.429)
Control mean	2.718	2.718	892.045	892.045
Observations	2148	2144	2183	2179
Controls	No	Yes	No	Yes

Notes: **Table A6 shows that jobseekers' beliefs about skills correlate positively with 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. Heteroskedasticity robust standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A7: Association between Skill Beliefs and Search Direction - Big Experiment Control Group

	Target numeracy jobs		Target communication jobs		Target concept formation jobs	
	(1)	(2)	(3)	(4)	(5)	(6)
Believed numeracy comp. adv.	0.100***	0.090***	-0.094***	-0.080***	-0.003	-0.008
	(0.014)	(0.014)	(0.023)	(0.022)	(0.015)	(0.014)
Believed communication comp. adv.	-0.079***	-0.071***	0.087***	0.080***	-0.000	0.002
	(0.021)	(0.020)	(0.024)	(0.025)	(0.022)	(0.022)
Believed concept formation comp. adv.	-0.013	-0.018	-0.029	-0.024	0.067***	0.064***
	(0.019)	(0.018)	(0.017)	(0.019)	(0.019)	(0.020)
Control mean	0.222	0.222	0.471	0.471	0.220	0.220
Observations	2183	2179	2183	2179	2183	2179
Controls	No	Yes	No	Yes	No	Yes

Notes: **Table A7 shows that jobseekers' beliefs about comparative advantage correlate positively with skill-directed search in the control group of the big experiments.** Dependent variables are dummies indicating that jobseekers rate the respective skill as most important for the jobs that they applied to in the last 30 days. Independent variables are dummies for beliefs indicating that individuals have a clear comparative advantage in each skill. Even columns control 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 in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.2 Appendix Exhibits for the Tight Experiment

This section contains tables that related to the tight, mechanism-focused field experiment. Table A8 shows treatment effects on beliefs about general skills in the tight experiment using the full sample of workshop participants. Table A9 displays the main treatment effects on aligned search direction outcomes using the full sample of workshop participants. Table A10 shows the balance tests for the tight experiment for the restricted and full sample. Table A11 shows robustness checks for the main results. Table A12 displays coefficients from regressions of an indicator for aligned comparative advantage belief with baseline demographics. Table A13 shows treatment effect heterogeneity using a non-binary measure of aligned comparative advantage beliefs. Table A14 shows treatment effect heterogeneity by baseline comparative advantage beliefs and confidence simultaneously.

We also elicit respondents' full discrete distribution of beliefs of skill quintiles for general skills. Figure A2 shows that the distribution of beliefs is closer to assessed skills for the treatment than the control group. Table A15 shows that participants also become somewhat more certain in their beliefs about general skills.[34] Average variance in beliefs goes down by 11%, with similar decreases across numeracy and communication beliefs, although this is noisily estimated.

In Table A16 we assess whether jobseekers learn about themselves or the distribution of skills in the reference population. To do so we ask jobseekers after the assessments but before treatment how many questions they correctly answered in the assessments. We then study how beliefs about absolute and relative performance predict belief updating. For comparative advantage beliefs, we find suggestive evidence that belief updating about the distribution is more important (columns 1 to 3). For beliefs about skill quintiles, we find suggestive evidence that jobseekers updating about their own level of skill is more important (columns 4 to 6).

Table A17 displays treatment effects on dummies indicating the fraction of aligned, under, and overconfident skill beliefs. Table A18 combines the tight and big experiment to show that there are no heterogeneous treatment effects on belief updating by the level of concept formation skill. Table A19 displays treatment effects on skill beliefs by gender. Table A20 shows summary statistics of participants' engagement with the job search platform (SAYouth.mobi) 30 days prior to the intervention in the tight experiment. Table A21 displays treatment effects on search direction in the tight experiment by whether skill requirements were revealed.

[34] Specifically, we ask respondents to allocate ten tokens across the five quintiles to represent the likelihood of them falling in that group. We use this data to calculate the expected skill quintile ($\bar{s}$) and calculate the variance in beliefs as $var(s_l) = \sum_{i=0}^{4} p_i(i - \bar{s})^2$. We then average the variance across the two skill domains.

Table A8: Treatment Effects on Skill Beliefs - Tight Experiment Including Jobseekers Without Unique Comparative Advantage

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

Notes: **Table A8 shows that treatment effects on beliefs about skills are robust to including people without a clear comparative advantage in skills.** Columns 1 and 2 show effects on a dummy indicating beliefs about respondents' comparative advantage in skills that are aligned with the assessment results. Columns 3 and 4 show treatment effects on the fraction of skill quintile beliefs that are aligned with the assessment results. Columns 2 and 4 show treatment effect heterogeneity by whether individuals had aligned comparative advantage beliefs at baseline. Controls include randomization block fixed effects, and prespecified baseline covariates (skill quintile dummies, age, gender, and dummies for having completed a university degree, other post-secondary education, and high school). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A9: Treatment Effects on Aligned Search Direction - Tight Experiment Including Jobseekers Without Unique Comparative Advantage

	Aligned search index		% aligned (job choice)		Δ SMS click rate		Δ planned apps (w)		Δ % aligned platform apps	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment	0.240***	0.435***	0.032	0.062**	0.071	0.122*	1.121	3.363***	0.052***	0.062**
	(0.078)	(0.094)	(0.025)	(0.026)	(0.046)	(0.063)	(0.878)	(1.140)	(0.017)	(0.030)
Treatment × Aligned comp adv belief (bl)		-0.445***		-0.070**		-0.113		-5.190**		-0.022
		(0.156)		(0.029)		(0.098)		(2.162)		(0.060)
Aligned comp adv belief (bl)		0.504***		0.111***		0.070		6.113***		0.010
		(0.121)		(0.026)		(0.076)		(1.872)		(0.035)
Treatment effect: Aligned comp adv belief (bl)		-0.009		-0.008		0.009		-1.827		0.041
		(0.118)		(0.029)		(0.070)		(1.613)		(0.039)
Control mean	0.000	0.000	0.416	0.416	-0.024	-0.024	3.272	3.272	0.005	0.005
Observations	368	368	368	368	368	368	368	368	368	368

Notes: **Table A9 shows that treatment effects on aligned search direction are robust to including people without a clear comparative advantage in skills.** The table includes the full sample of workshop participants. All search direction measures are coded as zero for jobseekers with tied skill quintiles. Aligned job search is defined as directing search effort toward jobs that mostly require the skill that aligns with job seekers' measured comparative advantage. Columns 1 and 2 show impacts on an index of search direction. The index is constructed as the variance-covariance weighted average of the search alignment measures displayed in columns 3 to 10 (Anderson, 2008). Columns 3 and 4 show the impact on the percentage of 11 incentivized job choices that are aligned with the comparative advantage. Columns 5 and 6 show the impact on the difference in link click rates between aligned and non-aligned jobs sent to job seekers via SMS. All job seekers receive three links with at least one job for each skill. Columns 7 and 8 show impact on the difference between aligned and non-aligned applications (winsorized). Columns 9 and 10 show impact on the difference between the percentage of aligned and non-aligned applications on the online job search platform SAyouth.mobi. Jobs are classified using a list of 20 numeracy-heavy and 20 communication-heavy common job titles classified by experts. Even columns show heterogeneity by whether individuals have aligned comparative advantage beliefs at baseline. Controls include randomization block fixed effects, and prespecified baseline covariates (skill quintile dummies, age, gender, dummies for having completed a university degree, other post-secondary certificate or diploma, and high school, and a dummy indicating a comparative advantage in numeracy). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A10: Balance Table - Tight Experiment

	Restricted sample					Full sample				
	(1) Control	(2) Treatment	(3) Δ	(4) p(Δ = 0)	(5) N	(6) Control	(7) Treatment	(8) Δ	(9) p(Δ = 0)	(10) N
Panel A: Demographics										
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	0.67	0.71	0.04	0.48	278	0.67	0.68	0.01	0.92	372
University degree / diploma	0.09	0.08	-0.01	0.91	278	0.10	0.06	-0.04	0.14	372
Other post-secondary education	0.14	0.14	-0.01	0.61	278	0.14	0.16	0.02	0.40	372
Panel B: Labor market background										
Any work in the 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 2021 USD (last 7 days, winsorized)	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
Panel C: Search behavior										
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, winsorized)	11.31	8.69	-2.62	0.10	278	11.46	10.49	-0.97	0.54	372
Search expenditure in 2021 USD (last 7 days, winsorized)	23.98	21.47	-2.51	0.25	278	24.42	22.65	-1.77	0.28	372
Hours spent searching (last 7 days, winsorized)	13.75	13.90	0.15	0.92	278	14.06	14.23	0.17	0.88	371
Job offers (last 30 days, winsorized)	0.14	0.21	0.07	0.11	278	0.20	0.23	0.03	0.51	372
Panel D: Search alignment with comparative advantage										
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	0.76	278	-0.00	-0.00	0.00	0.77	372
Panel E: Aligned comp. adv. belief										
General skill	0.47	0.50	0.04	0.68	278	0.43	0.46	0.03	0.69	369
Assessment result	0.53	0.54	0.01	0.95	278	0.45	0.46	0.02	0.82	371
Panel F: % aligned beliefs										
Assessment result	0.18	0.15	-0.03	0.34	278	0.15	0.14	-0.02	0.55	371
General skill	0.18	0.26	0.08	0.12	278	0.18	0.24	0.06	0.17	369

Notes: **Table A10 shows that covariates are balanced across treatment groups in the tight experiment.** Columns 1 to 5 show results for the sample of individuals with a clear comparative advantage in skills. Columns 6 to 10 show results for the full sample of individuals including those without unique comparative advantage in skills.

Table A11: Robustness Checks for Main Outcomes - Comparing Tight and Big Experiments

	Tight experiment		Big experiment			
	(1)	(2)	(3)	(4)	(5)	(6)
	Aligned comp. adv. belief	Aligned search index	Aligned comp. adv. belief	Aligned search direction	LM quantity index	LM quality index
Panel A: Robustness for inference						
Treatment	0.135	0.241	0.139	0.050	0.045	0.085
unadjusted p-value	(0.001)	(0.023)	(0.000)	(0.000)	(0.165)	(0.017)
wild-bootstrapped p-value	[0.002]	[0.015]	[0.000]	[0.000]	[0.390]	[0.004]
q-value	{0.002}	{0.014}	{0.001}	{0.001}	{0.029}	{0.013}
Panel B: Reweighting to match other sample						
Treatment	0.152***	0.241**	0.142***	0.050***	0.054*	0.089**
	(0.044)	(0.101)	(0.011)	(0.010)	(0.031)	(0.036)
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 A11 shows that the main results are robust to multiple hypothesis testing, using wild-bootstraps, and reweighting results to match the demographics across experiments.** Panel A shows unadjusted, wild-bootstrapped p-values, as well as sharpened q-values (Benjamini et al., 2006) estimated across all outcomes. wild-bootstrapped p-values are obtained using 10,000 draws. Panel B shows the same treatment effects using weights to match the observational characteristics of the other experimental sample. The covariates used to estimate weights are dummies for being male, just having matric, having a post-secondary degree, or having a post-secondary certificate, having worked in the last seven days, having searched for work in the last seven days, and for having ever worked in a wage job as well as the number of job applications in the last seven days. Columns 1 and 2 show effects on the main outcomes of the tight experiment. Columns 3 to 6 show effects on earnings conditional on doing any work. Controls are the same as in the main specifications. Standard errors clustered at the treatment-day level in parenthesis in panel B. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ in Panel B.

Table A12: Association Between Baseline Aligned Comparative Advantage Belief and Other Baseline Characteristics - Tight Experiment

	(1) Aligned CA belief	(2) Aligned CA belief
Age	-0.012	-0.004
	(0.009)	(0.008)
Female	0.048	0.024
	(0.064)	(0.059)
Has completed secondary education only	-0.008	-0.024
	(0.114)	(0.109)
Has a post secondary certificate	-0.059	-0.071
	(0.133)	(0.126)
Has a post secondary diploma	0.095	0.058
	(0.140)	(0.135)
Has a post secondary degree	-0.208	-0.208
	(0.140)	(0.140)
Employed in any form at baseline	0.030	0.026
	(0.068)	(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***	0.076
	(0.244)	(0.283)
# jobseekers	278	278
R2 (not adjusted)	0.031	0.177
p: all coefficients = 0	0.233	0.000

Notes: **Table A12 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. Heteroskedasticity-robust standard errors shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A13: Heterogeneous Treatment Effects by Non-binary Baseline Comparative Advantage Beliefs - Tight Experiment

	Main outcomes				Additional belief outcomes			
	Aligned search index		Aligned comp. adv. belief		Degree of CA alignment		Overall alignment of belief levels	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	0.255**	0.856**	0.152***	0.444***	0.063***	0.324***	0.073***	0.138**
	(0.101)	(0.396)	(0.044)	(0.142)	(0.012)	(0.091)	(0.012)	(0.062)
Treatment × Baseline degree of CA belief alignment		-0.721		-0.348*		-0.315***		-0.078
		(0.486)		(0.176)		(0.106)		(0.072)
Baseline degree of CA belief alignment		1.274**		1.125***		0.828***		0.045
		(0.496)		(0.215)		(0.071)		(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 A13 shows that treatment effect heterogeneity and treatment effects on beliefs are robust to using a non-binary measure of aligned comparative advantage beliefs in the tight experiment.** The non-binary measure of comparative advantage beliefs captures the degree to which comparative advantage beliefs are misaligned and ranges from 0 to 1. If comparative advantage beliefs are aligned it is set to 1. If they are misaligned defined as $1 - abs(\Delta bel_skill_{i,j} + 1)/5$, where $\Delta bel_skill_{i,j}$ is the difference between numeracy and communication beliefs (in quintiles). A value of zero means that jobseekers need to shift their skill beliefs by a total of five quintiles to align their comparative advantage belief with the measured comparative advantage. The mean of this measure is 0.83 and the standard deviation is 0.22. Columns 1 to 4 show robustness of the main result heterogeneity to using a non-binary measure of aligned comparative advantage beliefs. Columns 1 and 2 show effects on the aligned search direction index described in table 6. Columns 3 and 4 show effects on a dummy indicating aligned comparative advantage beliefs. Columns 5 to 8 show treatment effects on additional measures of belief alignment. Columns 5 and 6 show treatment effects on a non-binary measure of aligned comparative advantage beliefs. Columns 7 and 8 show treatment effects on a measure of overall alignment of skill beliefs with measured skills. This measure is defined as one minus the average difference between beliefs about skill quintiles and assessed quintiles relative to maximum possible difference. A value of 1 indicates perfectly aligned beliefs while a value of zero indicates maximally misaligned beliefs. Even columns show heterogeneity by the baseline value of the non-binary measure of comparative advantage beliefs. Controls include randomization block fixed effects, and prespecified baseline covariates (skill quintile dummies, age, gender, and dummies for having completed a university degree, other post-secondary education, and high school). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

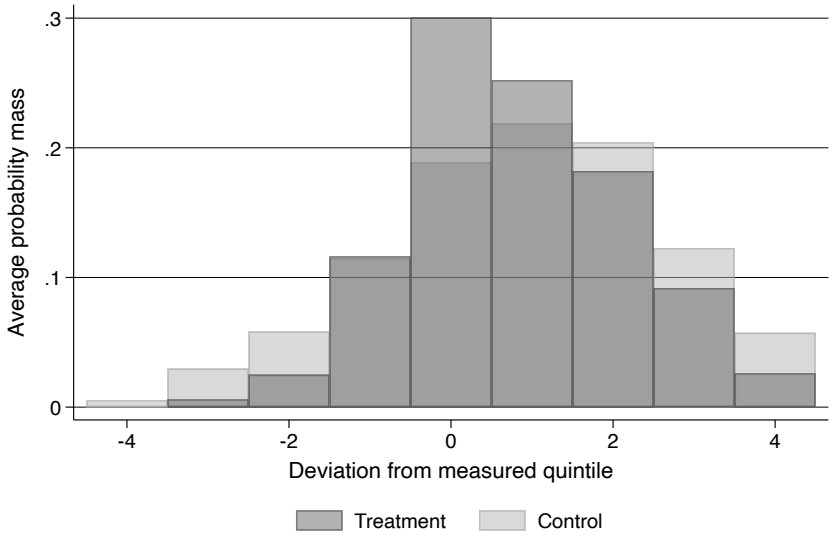
Table A14: Heterogeneous Treatment Effects on Main Outcomes by Comparative Advantage Beliefs and Confidence - Tight Experiment

	(1) Aligned comp. adv. belief	(2) Aligned search index	(3) Search effort index
Treatment	0.216***	0.551***	-0.194
	(0.066)	(0.155)	(0.177)
Treatment × Aligned comp adv belief (bl)	-0.161	-0.651**	0.276
	(0.112)	(0.247)	(0.298)
Aligned comp adv belief (bl)	0.737***	0.539**	-0.056
	(0.110)	(0.212)	(0.226)
Treatment × Average confidence (bl)	-0.003	0.084	0.082
	(0.079)	(0.179)	(0.176)
Average confidence (bl)	0.072	0.061	0.200
	(0.058)	(0.142)	(0.143)
Treatment × Aligned comp adv belief (bl) × Average confidence (bl)	0.013	-0.064	-0.170
	(0.111)	(0.208)	(0.230)
Control mean	0.475	0.000	0.000
Observations	278	278	278

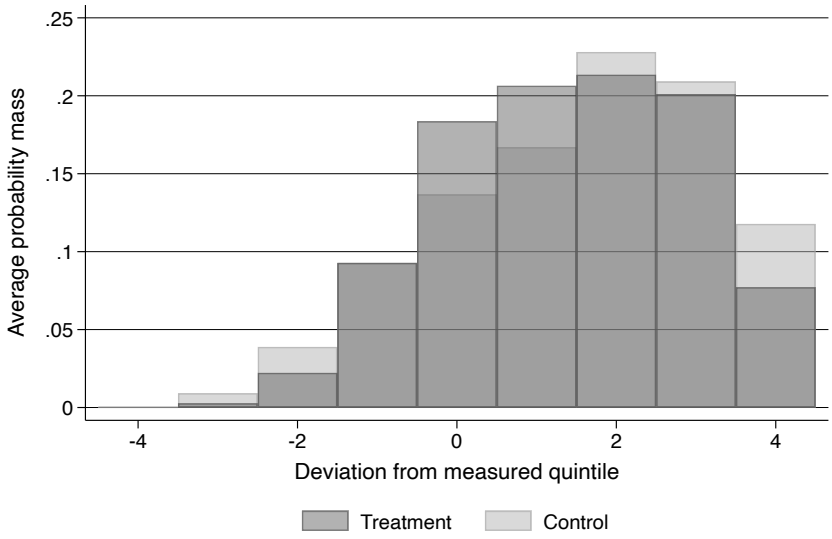
Notes: **Table A14 shows that the main results are robust to interacting the treatment with both baseline comparative advantage beliefs and baseline confidence in skills.** The magnitude of the interaction term between the treatment dummy and baseline comparative advantage does not change much relative to Tables 4 and 6. Moreover, the magnitude of the three-way interaction term is small and insignificant throughout suggesting that the confidence levels do not affect treatment effect heterogeneity by baseline comparative advantage beliefs. Baseline confidence levels are defined as the average baseline deviation of quintile skill beliefs from measured quintile (across numeracy and communication) standardized to have control group standard deviation of one. Positive values indicate overconfidence and negative values indicate underconfidence. Column 1 shows effects on a dummy indicating aligned comparative advantage beliefs. Column 2 shows effects on the aligned search index. Column 3 shows effects on the search effort index. Controls are the same as in the main specifications. Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A2: Treatment Effects on Distribution of Beliefs About Skills - Tight Experiment

Panel A: Numeracy beliefs



Panel B: Communication beliefs



Notes: **Figure A2 shows how the treatment shifted beliefs about skills toward assessed skills in the tight experiment.** Beliefs are elicited by asking respondents to place 10 buttons across the five possible quintiles. x-axis measures the deviation from the assessed quintile. y-axis measures the average probability mass. Panel A displays the distribution of beliefs for numeracy beliefs. Panel B displays the distribution of communication beliefs.

Table A15: Treatment Effects on The Variance of Beliefs About Skills - Tight Experiment

	Average variance in beliefs		Variance numeracy beliefs		Variance communication beliefs	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-0.049	-0.071	-0.029	-0.062	-0.068	-0.080
	(0.054)	(0.077)	(0.058)	(0.087)	(0.055)	(0.076)
Treatment × Aligned comp adv belief (bl)		0.056		0.077		0.036
		(0.129)		(0.147)		(0.129)
Aligned comp adv belief (bl)		-0.148		-0.145		-0.152
		(0.095)		(0.101)		(0.099)
Treatment effect:		-0.015		0.014		-0.045
Aligned comp adv belief (bl)		(0.089)		(0.098)		(0.092)
Control mean	0.644	0.644	0.679	0.679	0.609	0.609
Observations	278	278	278	278	278	278

Notes: **Table A15 shows treatment effects on the variance of general, non-assessment specific, skill beliefs in the tight experiment.** Domain-specific beliefs about skill quintiles were elicited using the histogram method. Column 1 shows effects on the average variance. Column 2 shows effects on the variance of numeracy beliefs. Column 3 shows effects on the variance of communication skills. All specifications include randomization block fixed effects. Controls further include prespecified baseline covariates (skill quintile dummies, age, and dummies for having completed a university degree, other post-secondary education, and high school). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A16: Learning About Oneself or The Distribution of Skills?

	Aligned comp. adv. belief			Av. abs. deviation from assessments		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.171**	0.208***	0.176***	-0.090***	-0.086***	-0.090***
	(0.065)	(0.050)	(0.060)	(0.014)	(0.014)	(0.015)
Treatment × Aligned abs. comp adv.	-0.037		0.139			
	(0.096)		(0.150)			
Aligned abs. comp adv.	0.473***		0.117			
	(0.118)		(0.138)			
Treatment × Aligned rel. comp adv.		-0.137	-0.240			
		(0.082)	(0.148)			
Aligned rel. comp adv.		0.586***	0.543***			
		(0.079)	(0.112)			
Treatment × Abs. dev. % correct (z)				-0.030**		-0.029**
				(0.013)		(0.013)
Abs. dev. % correct (z)				0.032***		0.025*
				(0.011)		(0.013)
Treatment × Abs. dev. quintile (z)					0.012	0.006
					(0.013)	(0.012)
Abs. dev. quintile (z)					-0.022**	-0.014
					(0.009)	(0.009)
Control mean	0.475	0.475	0.475	0.360	0.360	0.360
Observations	278	278	278	278	278	278

Notes: **Table A16 explores whether belief updating is driven by learning about oneself or about the distribution of skills by others.** Columns 1 to 3 show effects on beliefs about jobseekers' comparative advantage. Columns 4 to 6 show effects on the expected absolute deviation of beliefs about skills (elicited using the histogram method) from measured quintiles. Columns 1 and 4 display heterogeneity by baseline measures defined using beliefs about the absolute performance on the assessments (in percent correctly answered questions). Columns 2 and 5 display heterogeneity by baseline measures defined using the relative performance on assessments (in quintiles). Columns 3 and 6 display heterogeneity by both measures in the same regression. All specifications include randomization block fixed effects. Controls further include prespecified baseline covariates (skill quintile dummies, age, and dummies for having completed a university degree, other post-secondary education, and high school). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A17: Treatment Effects on Over- and Underconfident Beliefs - Tight Experiment

	Framed field experiment (quintiles)		Field experiment (terciles)	
	(1) Underconfident	(2) Overconfident	(3) Underconfident	(4) Overconfident
Treatment	-0.057***	-0.022	-0.042***	-0.101***
	(0.016)	(0.018)	(0.005)	(0.007)
Treatment effect/ control mean	-0.306	-0.035	-0.282	-0.219
	(0.087)	(0.028)	(0.033)	(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 A17 shows that underconfident beliefs are more likely to be updated than overconfident beliefs.** It displays treatment effects on dummies indicating the fraction of aligned, under, and overconfident skill beliefs. Columns 1 to 2 show results for the tight experiment. Columns 3 and 4 show results of the field experiment. Columns 1 and 3 show effects on a dummy indicating the fraction of skill domains on which beliefs about respondents' skills are underconfident relative to the assessment results. Columns 2 and 4 show effects on a dummy indicating the fraction of skill domains on which beliefs about respondents' skills are overconfident relative to the assessment results. The effect sizes relative to the mean suggest that underconfident beliefs are more likely to update. All specifications include randomization block fixed effects. Controls further include prespecified baseline covariates. Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A18: Heterogeneous Treatment Effects on Belief Updating by Cognitive Skill - Both Experiments

	Fraction aligned	
	(1) Tight experiment	(2) Big experiment
Treatment	0.079	0.164
	(0.026)	(0.011)
Treatment × CFT	0.018	0.005
	(0.020)	(0.008)
CFT	-0.004	0.005
	(0.016)	(0.007)
Control mean	0.183	0.410
Observations	278	4205
Number of clusters	34	54

Notes: **Table A18 shows that treatment effects on belief updating do not vary by a proxy for cognitive skill.** The proxy is the standardized score on the concept formation assessment, which is similar to the Raven's matrices assessment. Column 1 regresses the share of aligned beliefs on treatment, standardized CFT score (Raven's matrices), their interaction and the usual controls in the tight experiment. Column 2 does the same for the big experiment, using the standardized concept formation test score and the respective controls. Standard errors clustered at the treatment-day level.

Table A19: Heterogeneous Treatment Effects on Skill Beliefs by Gender - Tight Experiment

	Aligned com. adv. 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 × 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	278	278	278
Controls	No	Yes	No	Yes

Notes: **Table A19 shows that treatment effects on skill beliefs do not differ by gender.** Columns 1 and 2 show effects on a dummy indicating beliefs about respondents' comparative advantage in skills that are aligned with the assessment results. Columns 3 and 4 show treatment effects on the fraction of skill beliefs that align with measured skill quintiles. All specifications include randomization block fixed effects. Controls further include prespecified baseline covariates (skill quintile dummies, age, and dummies for having completed a university degree, other post-secondary education, and high school). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A20: Summary Statistics for Search on SAYouth.mobi Platform - Tight Experiment

	mean	sd	median	min	max	N
# days active on platform	3.17	4.06	2.00	0.00	25.00	278
# applications clicks (w)	6.51	11.21	2.00	0.00	65.00	278
# applications clicks for numeracy heavy jobs	0.39	1.10	0.00	0.00	11.00	278
# applications clicks for communication heavy jobs	0.82	1.64	0.00	0.00	9.00	278
Fraction of skill coded application clicks	0.12	0.20	0.00	0.00	1.00	278

Notes: **Table A20 shows summary statistics of participants' engagement with the job search platform (SAYouth.mobi) 30 days prior to the intervention in the tight experiment.** Application clicks are defined as initiating an application by clicking on the button "apply here".

Table A21: Heterogeneous Treatment Effects on Search Direction for Jobs With and Without Revealed Skill Demands - Tight Experiment

	Choice aligned with comparative advantage	
	(1) Misaligned baseline beliefs	(2) Aligned baseline beliefs
Treatment	0.039	0.023
	(0.032)	(0.034)
Treatment × Skill req. revealed	0.130**	-0.114
	(0.063)	(0.085)
Treatment effect: Skill req. revealed	0.169***	-0.091
	(0.060)	(0.077)
Control mean	0.413	0.413
Observations	1573	1485

Notes: **Table A21 shows that, for jobseekers with initially misaligned comparative advantage beliefs, treatment effects on search direction in the job choice tasks are stronger when skill requirements are revealed.** The table shows the impact on dummy variables indicating job choices aligned with jobseekers' assessed comparative advantage estimated at the job-choice individual level. Column 1 shows the impact for individuals with initially misaligned beliefs. Column 2 shows the impact for individuals with initially aligned beliefs. The sample is restricted to individuals with a clear comparative advantage in skills. Controls include randomization block fixed effects, job pair and job pair order fixed effects, and pre-specified baseline covariates (skill quintile dummies, age, gender, dummies for having completed a university degree, other post-secondary certificate or diploma, and high school, and a dummy indicating a comparative advantage in numeracy). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.3 Appendix Tables for the Big Experiment

This section contains additional tables and figures that relate to the big field experiment discussed in Section 4. Table A22 shows summary statistics for the big experiment. Figure A3 shows the skill report used in the private treatment. Table A23 shows that treatment did not induce differential attrition. Tables A24 shows balance tests for the full and non-attritted sample, respectively. Table A25 shows heterogeneous treatment effects on skill beliefs by respondents' gender. Figure A4 shows the cumulative distribution function of earnings in the last seven days in the field experiments. Table A26 shows treatment effects on selected transformations of earnings in the last seven days. Table A27 shows treatment effects on selected work quality outcomes multiplied by potential mediators. Table A28 shows that heterogeneous treatment effects on skill beliefs and search direction by baseline aligned comparative advantage belief are consistent with the tight experiment. Table A29 shows that heterogeneous treatment effects on labor market outcomes by aligned baseline comparative advantage beliefs in the big experiment are consistent with the tight experiment.

Table A22: Summary Statistics - Big Experiment

	(1) Mean	(2) Median	(3) SD	(4) Min	(5) Max	(6) Obs.
Panel A: Demographics						
Black African	0.92	1.00	0.27	0.00	1.00	4389
Male	0.38	0.00	0.48	0.00	1.00	4389
Age	23.67	23.14	3.28	18.04	35.08	4389
Completed secondary education only	0.61	1.00	0.49	0.00	1.00	4389
Any other post-secondary qualification	0.22	0.00	0.41	0.00	1.00	4389
University degree / diploma	0.17	0.00	0.37	0.00	1.00	4389
Panel B: Labor market background						
Worked in past 7 days	0.37	0.00	0.48	0.00	1.00	4389
Earnings in USD (last 7 days, winsorized)	31.26	0.00	75.72	0.00	476.00	4389
Has worked in permanent wage job before	0.09	0.00	0.29	0.00	1.00	4377
Panel C: Search behavior						
Did any search in past 7 days	0.97	1.00	0.17	0.00	1.00	4389
# applications (last 30 days, winsorized)	9.34	5.00	12.85	0.00	90.00	4346
Search expenditure in USD (last 7 days, winsorized)	30.97	20.40	33.05	0.00	204.00	3995
# job offers (last 30 days, winsorized)	0.80	0.00	2.66	0.00	20.00	4335
Panel D: Skills beliefs						
Aligned belief about comparative advantage	0.20	0.00	0.40	0.00	1.00	4312
Fraction of aligned skill beliefs	0.38	0.33	0.31	0.00	1.00	4378
beliefs and assessments						

Notes: **Table A22 shows summary statistics for the big experiment.**

Figure A3: 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.

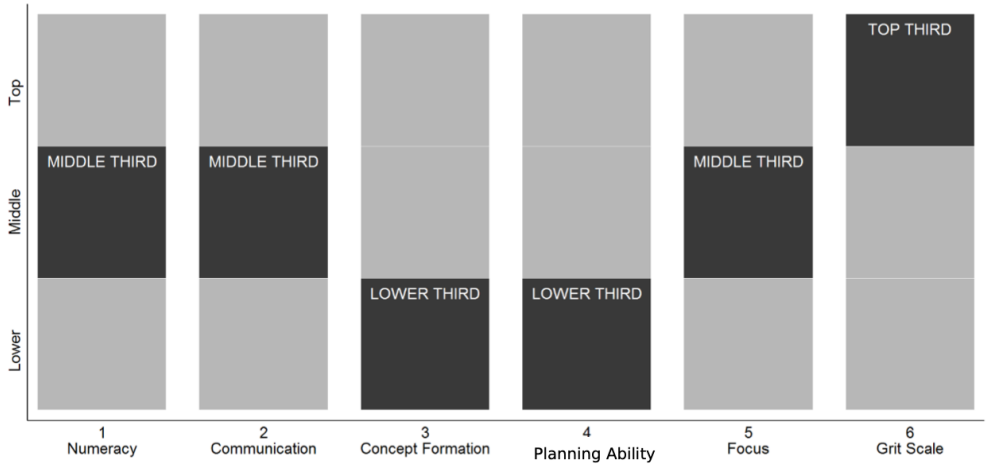
1. The Numeracy tests measure various maths abilities. Your score is the average of the two maths tests you did today at Harambee.
2. The Communication test measures English language ability through listening, reading and comprehension.
3. The Concept Formation test measures the ability to understand 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.

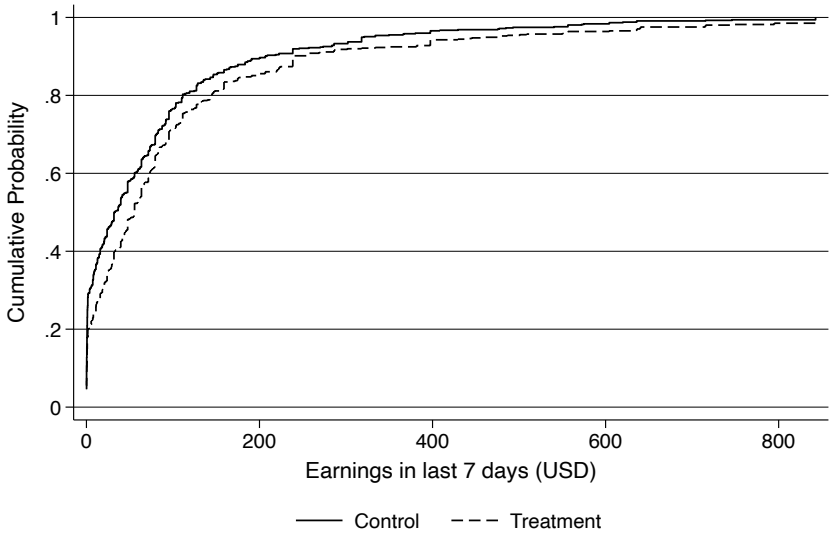


DISCLAIMER
Please note that this is a confidential assessment report and is intended for use by the person specified above. Assessment results are not infallible and may not be entirely accurate.

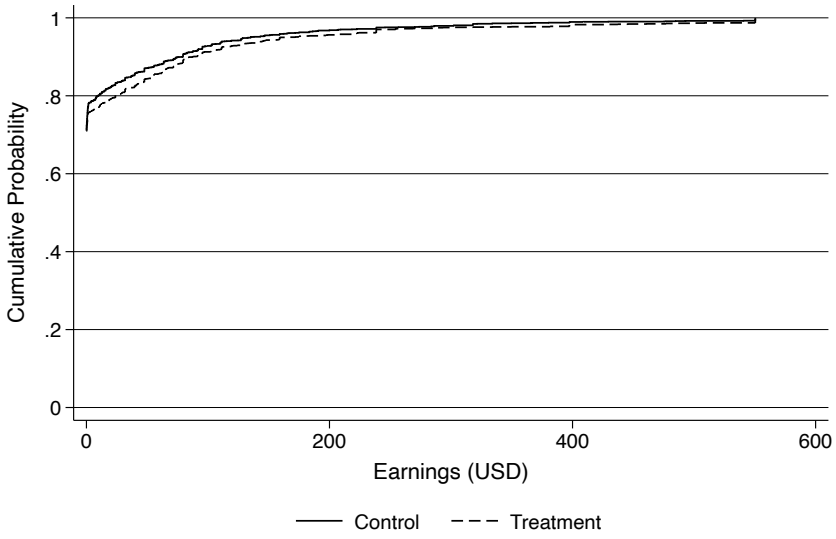
Notes: **Figure A3 shows an example of the reports given to treated jobseekers in the big experiment.** Each report contains the jobseeker's assessment results but no identifying information and no branding.

Figure A4: Cumulative Distributions of Earnings by Treatment Group - Big Experiment

Panel A: Conditional earnings



Panel B: Unconditional earnings



Notes: **Figure A4 shows that the earnings effects in the big experiment are not driven by outliers.** It shows the cumulative distribution function of earnings in the last seven days in the field experiments. Earnings measured in 2021 USD PPP and winsorized at the 99^{th} percentile. Panel A shows earnings conditional on having worked in the last seven days. Panel B shows unconditional earnings.

Table A23: Treatment Effects on Completing Endline Survey - Big Experiment

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

Notes: **Table A23 shows that attrition is low and balanced across treatment groups in the big experiment.** Dependent variable is a dummy indicating whether an individual was recontacted at for the endline survey. Column 2 includes pre-specified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, ***

Table A24: Balance Table - Big Experiment

	Full sample					Non-attritted sample				
	(1) Control	(2) Treatment	(3) Δ	(4) p(Δ = 0)	(5) N	(6) Control	(7) Treatment	(8) Δ	(9) p(Δ = 0)	(10) N
Panel A: Demographics										
Black African	0.93	0.91	-0.03	0.34	4389	0.93	0.91	-0.03	0.30	4206
Male	0.39	0.36	-0.03	0.04	4389	0.39	0.36	-0.03	0.03	4206
Age	23.55	23.79	0.29	0.07	4389	23.53	23.80	0.31	0.05	4206
Completed secondary education only	0.62	0.59	-0.04	0.20	4389	0.62	0.59	-0.04	0.17	4206
Any other post-secondary qualification	0.21	0.22	0.01	0.61	4389	0.22	0.23	0.01	0.63	4206
University degree / diploma	0.16	0.18	0.03	0.29	4389	0.15	0.18	0.04	0.24	4206
Panel B: Labor market background										
Worked in past 7 days	0.36	0.39	0.02	0.22	4389	0.36	0.39	0.02	0.19	4206
Earnings in USD (last 7 days, winsorized)	30.09	32.52	3.58	0.13	4389	29.84	32.49	3.79	0.11	4206
Has worked in permanent wage job before	0.09	0.09	-0.01	0.61	4377	0.10	0.09	-0.01	0.35	4195
Panel C: Search behavior										
Did any search in past 7 days	0.97	0.97	0.01	0.07	4389	0.97	0.98	0.01	0.02	4206
# applications (last 30 days, winsorized)	9.31	9.37	0.03	0.95	4346	9.13	9.27	0.13	0.79	4165
Search expenditure in USD (last 7 days, winsorized)	32.09	31.14	-1.11	0.34	3912	32.02	31.15	-0.95	0.42	3747
# job offers (last 30 days, winsorized)	0.77	0.84	0.04	0.64	4335	0.78	0.87	0.07	0.44	4152
Panel D: Skills beliefs										
Aligned belief about comparative advantage	0.20	0.21	0.01	0.49	4312	0.20	0.21	0.01	0.67	4132
% aligned beliefs	0.38	0.38	-0.01	0.49	4378	0.38	0.37	-0.01	0.45	4196

Notes: **Table A24 shows that covariates are balanced across treatment arms for the big experiment.** Columns 1 to 5 show statistics for the baseline sample. Columns 6 to 10 show statistics for the sample of individuals reached at endline.

Table A25: Heterogeneous Treatment Effects on Beliefs about Skills by Gender - Big Experiment

	Aligned comp. adv. 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 × 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		X		X

Notes: **Table A25 shows that treatment effects on skill beliefs do not differ by gender in the big experiment.** Columns 1 and 2 show impacts on a dummy indicating comparative advantage beliefs that align with assessment results. Columns 3 and 4 show impacts on the fraction of skill beliefs that align with measured skill terciles across three skill domains. Even columns include prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A26: Treatment Effects on Different Transformations of Earnings - Big Experiment

	Unconditional earnings					Conditional earnings				
	(1) raw	(2) wins	(3) z	(4) ln(earn+1)	(5) ihs	(6) raw	(7) wins	(8) z	(9) ln(earn+1)	(10) ihs
Treatment	9.360**	6.517**	0.097**	0.115**	0.127**	25.100**	20.393**	0.159**	0.229**	0.244**
	(3.610)	(2.712)	(0.037)	(0.053)	(0.060)	(11.348)	(9.404)	(0.072)	(0.093)	(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 A26 shows that treatment effects on earnings in the last seven days are robust to a range of functional forms.** Earnings are measured in 2021 USD PPP. Columns 1 to 5 show effects on unconditional earnings. Columns 6 to 10 show effects on earnings conditional on doing any work. Columns 1 and 6 show effects on the raw variable. Columns 2 and 7 show effects on the variable winsorized at the 99^{th} percentile. Columns 3 and 8 show effects on the variable standardized using the control group mean and standard deviation. Columns 4 and 9 show effects on the natural logarithm of the variable plus one. Columns 5 and 10 show effects on the inverse hyperbolic sine of the variable. Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A27: Treatment Effects on Combinations of Labor Market Outcome Measures - Big Experiment

	Earnings (w)					Written contract	
	(1) Start before treatment	(2) Start after treatment	(3) Tenure	(4) Wage employment	(5) Self employment	(6) Start before treatment	(7) Start after treatment
Treatment	1.985*	4.607**	10.303	6.923***	-0.405	0.003	0.014
	(1.072)	(2.237)	(8.297)	(2.562)	(0.694)	(0.004)	(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 A27 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 multiplied by potential mediators. Columns 1 to 5 show impacts on earnings in the last seven days (winsorized at the 99^{th} percentile). Columns 6 and 7 show impacts on a written contract dummy. Column titles indicate the variable with which the outcome is multiplied. Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). Standard errors clustered at the treatment-day level in parenthesis. All monetary figures are reported in 2021 USD PPP. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A28: Heterogeneous Treatment Effects on Skill Beliefs and Search Direction by Aligned Baseline Comparative Advantage Belief - Big Experiment

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

Notes: **Table A28 shows that treatment effects by baseline aligned comparative advantage beliefs on skill beliefs and search direction in the big experiment are broadly in line with the results of the tight experiment.** Column 1 shows the impact on aligned comparative advantage beliefs. Column 2 shows the impact on the fraction of aligned belief domains. Column 3 shows the impact on a dummy indicating search direction aligned with comparative advantage in skills (self-reported). Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A29: Heterogeneous Treatment Effects on Labor Market Outcomes by Aligned Baseline Comparative Advantage Beliefs - Big Experiment

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

Notes: **Table A29 shows that treatment effects on labor market outcomes in the big experiment are driven by jobseekers with initially misaligned comparative advantage beliefs in line with the tight experiment.** Columns 1 to 5 show impacts on employment quantity. Column 1 shows impacts on an Anderson (2008) index of the four employment quantity measures. Column 2 shows the impact on the number of job offers in the last 30 days (winsorized). Column 3 shows the impact on a dummy indicating any work for pay in month 1 after treatment. Column 4 shows the impact on a dummy indicating any work for pay in month 2 after treatment. Column 5 shows the impact on a dummy indicating any work for pay in the last seven days. Columns 6 to 9 show impacts on employment quality. Column 6 shows impacts on an Anderson (2008) index of the three employment quality measures. Column 7 shows the impact on earnings in the last seven days (winsorized). Column 8 shows the impact on hourly wages in the last seven days (winsorized). Column 9 shows the impact on a dummy indicating a written contract. Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). All monetary figures are reported in 2021 USD PPP. Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.4 Additional Mechanism Tables

This appendix contains tables with an analysis of potential mechanisms other than comparative advantage beliefs discussed in Section 5. First, we explore the role of belief levels and search effort in Section 5.1. Table A30 shows impacts on belief levels. Tables A31, A31, and A32 show impacts of different measures of search effort. All five tables show both average treatment effects and heterogeneity by baseline confidence (a measure of skill beliefs relative to assessed skills). Table A34 shows heterogeneity in labor market outcomes by baseline confidence in skill beliefs.

We view the treatment effects on search effort as relatively similar across the two experiments and we cannot reject the equality of effects on the two indices (p = 0.47). However, some readers may wonder why one experiment produces a near-zero effect and the other produces a weakly negative effect. For these readers, we evaluate three possible explanations for this pattern. First, **measurement**: treatment effects might differ because the experiments use different search effort measures, designed to take advantage of their different designs. To test this explanation, we compare treatment effects on the most comparable measures across the two experiments: applications submitted on the platform for the tight experiment and applications reported in surveys for the big experiment. Treatment effects on these measures show similar patterns to treatment effects on the indices, providing some evidence against the measurement explanation (Tables A32 and A33). Second, **timing**: effects might be marginally negative soon after treatment, when the tight experiment measures are observed, and then move toward zero after several months when the big experiment measures are observed. To test this explanation, we estimate treatment effects on search effort in the big experiment by the time between treatment and the endline survey, instrumented by the randomized assignment to earlier or later survey attempts. Treatment effects do not go from negative to zero, providing some evidence of the timing explanation. Third, **dimensionality**: the tight experiment provides jobseekers with results on fewer assessments, so more jobseekers get negative signals for every skill than in the big experiment. However, we see no evidence that jobseekers who get more negative signals have different treatment effects on search effort, providing some evidence against the dimensionality explanation. We conclude that none of these three are probable explanations for a difference in search effort effects between the two experiments but emphasize that the effects are similar enough that there is little to explain.

The remaining tables in this appendix show results related to the other mechanisms discussed in Section 5. Table A35 shows evidence for two additional potential mechanisms in the big experiment: self-esteem and education investments. The latter are discussed in Section 5.3 jointly with willingness-to-pay results presented in Appendix Section C. Table A36 shows that average treatment effects on labor market outcomes are driven by jobseekers who attached their skill report with applications.

Table A30: Heterogeneous Treatment Effects on Skill Beliefs by Baseline Confidence

	Tight experiment				Big experiment			
	Average abs. deviation		Average skill quintile beliefs (0-4)		Average abs. deviation		Average skill tercile beliefs (0-2)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	-0.085***	-0.080***	-0.082**	0.157*	-0.192***	-0.157***	-0.109***	-0.007
	(0.014)	(0.017)	(0.035)	(0.083)	(0.012)	(0.015)	(0.010)	(0.011)
Treatment × Average confidence (bl)		-0.007		-0.263***		-0.043***		-0.125***
		(0.014)		(0.060)		(0.012)		(0.012)
Average confidence (bl)		0.021		0.512***		0.101***		-0.063***
		(0.016)		(0.071)		(0.019)		(0.017)
Control mean	0.360	0.360	2.693	2.693	0.734	0.734	1.446	1.446
Observations	278	278	278	278	4191	4118	4195	4131

Notes: **Table A30 shows that beliefs about relative skill level update downwards and differentially depending on jobseekers' baseline confidence about their relative skills.** Baseline confidence levels are defined as the average baseline deviation of quintile skill beliefs from the assessed quintile across numeracy and communication (tight experiment) or tercile across numeracy, communication, and concept formation (big experiment). Confidence measures are standardized to have a control group standard deviation of one. Positive values indicate overconfidence and negative values indicate underconfidence. Columns 1 to 4 show results for the tight experiment. Columns 5 to 8 show results for the big experiment. Columns 1, 2, 5, and 6 show impacts on the average absolute deviation of beliefs from measured skills. Columns 3, 4, 7, and 8 show impacts on average skill beliefs. Columns 2, 4, 6, and 8 show effects by baseline confidence. Controls include randomization block fixed effects, and prespecified baseline covariates. Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A31: Heterogeneous Treatment Effects on Search Effort - Tight Experiment

	Index		Planned apps (w)		Time spent drafting (w)		SMS click rate		Observed apps (w)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment	-0.078	-0.101	-3.854	-3.574	-0.530	-0.377	0.003	-0.055	-1.463	-2.058
	(0.103)	(0.144)	(2.555)	(4.124)	(0.591)	(0.925)	(0.032)	(0.053)	(1.553)	(2.385)
Treatment × Average confidence (bl)		0.017		-0.429		-0.205		0.062		0.576
		(0.111)		(2.936)		(0.771)		(0.045)		(2.073)
Average confidence (bl)		0.197		4.032		1.280		-0.059		1.239
		(0.124)		(2.675)		(0.900)		(0.046)		(1.918)
Control mean	-0.000	-0.000	37.878	37.878	8.828	8.828	0.635	0.635	15.216	15.216
Observations	278	278	278	278	267	267	278	278	278	278

Notes: **Table A31 shows that treatment effects on search effort in the tight experiment do not vary significantly by baseline confidence levels.** Baseline confidence levels are defined as the average baseline deviation of quintile skill beliefs from measured quintiles (across numeracy and communication) standardized to have control group standard deviation of one. Positive values indicate overconfidence and negative values indicate underconfidence. Columns 1 and 2 show impacts on an Anderson (2008) index of four search effort measures. Columns 3 and 4 show impacts on the number of planned applications in the next 30 days (winsorized). Columns 5 and 6 show impacts on the number of minutes individuals spent drafting a cover letter during the job-search workshop (winsorized). Columns 7 and 8 show impacts on the click rate for three SMS with links to job adverts sent to individuals. Columns 9 and 10 show impacts on the number of applications sent on the job-search platform in the 30 days following the treatment (winsorized). Controls include randomization block fixed effects, and prespecified baseline covariates. Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A32: 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	-0.238	-0.723	-1.451**	-0.206	-2.906	-1.463	-2.058
	(0.092)	(0.156)	(0.477)	(0.598)	(1.519)	(3.119)	(1.553)	(2.385)
Treatment × Average confidence (bl)		0.142		0.759		2.848		0.576
		(0.121)		(0.449)		(2.580)		(2.073)
Average confidence (bl)		-0.066		-0.278		-2.303		1.239
		(0.111)		(0.530)		(2.279)		(1.918)
Control mean	-0.000	-0.000	6.014	6.014	9.058	9.058	15.216	15.216
Observations	278	278	278	278	278	278	278	278

Notes: **Table A32 shows that treatment effects on prespecified search effort measures on the job-search platform SAYouth.mobi in the tight experiment do not vary significantly by baseline confidence levels.** Baseline confidence levels are defined as the average baseline deviation of quintile skill beliefs from measured quintile (across numeracy and communication) standardized to have control group standard deviation of one. Positive values indicate overconfidence and negative values indicate underconfidence. Columns 1 and 2 show impacts on an Anderson (2008) index of the three search effort measures. Columns 3 and 4 show impacts on the number of days jobseekers were active on the platform in the 30 days following the treatment. Columns 5 and 6 on the number of job adverts jobseekers clicked on on the platform in the 30 days following the treatment (winsorized). Columns 7 and 8 show impacts on the number of applications sent on the job-search platform in the 30 days following the treatment (winsorized). Controls include randomization block fixed effects, and prespecified baseline covariates. Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A33: 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	-0.018	0.056	0.010	-0.321	-0.601	0.243	-0.077
	(0.032)	(0.038)	(0.488)	(0.572)	(0.318)	(0.498)	(0.818)	(0.855)
Treatment × Average confidence (bl)		0.030		0.141		0.355		0.493
		(0.036)		(0.464)		(0.490)		(0.675)
Average confidence (bl)		-0.034		-0.073		-0.604		-0.416
		(0.049)		(0.679)		(0.657)		(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 A33 shows that treatment effects on search effort in the big experiment do not vary significantly by baseline confidence levels.** Baseline confidence levels are defined as the average baseline deviation of tercile skill beliefs from measured terciles (across numeracy, communication, and concept formation) standardized to have control group standard deviation of one. Positive values indicate overconfidence and negative values indicate underconfidence. Columns 1 and 2 show impacts on an Anderson (2008) index of three search effort measures. Columns 3 and 4 show impacts on the number of applications in the last 30 days (winsorized). Columns 5 and 6 show impacts on the number of hours spent searching for jobs in the last 30 days (winsorized). Columns 7 and 8 show impacts on job search expenditure in the last 30 days (winsorized). Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). All monetary figures are reported in 2021 USD PPP. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A34: Heterogeneous Treatment Effects on Labor Market Outcomes by Confidence - Big Experiment

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

Notes: **Table A34 shows that treatment effects on labor market outcomes in the big experiment do not vary by baseline confidence levels.** Baseline confidence levels are defined as the average baseline deviation of tercile skill beliefs from measured terciles (across numeracy, communication, and concept formation) standardized to have control group standard deviation of one. Positive values indicate overconfidence and negative values indicate underconfidence. Columns 1 to 5 show impacts on employment quantity. Column 1 shows impacts on an Anderson (2008) index of the four employment quantity measures. Column 2 shows the impact on the number of job offers in the last 30 days (winsorized). Column 3 shows the impact of a dummy indicating any work for pay in month 1 after treatment. Column 4 shows the impact of a dummy indicating any work for pay in month 2 after treatment. Column 5 shows the impact on a dummy indicating any work for pay in the last seven days. Columns 6 to 9 show impacts on employment quality. Column 6 shows the impact on an Anderson (2008) index of the three employment quality measures. Column 7 shows the impact on earnings in the last seven days (winsorized). Column 8 shows the impact on hourly wages in the last seven days (winsorized). Column 9 shows the impact on a dummy indicating a written contract. Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). All monetary figures are reported in 2021 USD PPP. Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A35: Treatment Effects on Additional Mechanisms - Big Experiment

	Self-esteem				Education Investment		
	(1) SMS (z)	(2) SMS above med.	(3) Endline (z)	(4) Endline above med.	(5) Any	(6) Apprenticeship	(7) Formal
Treatment	0.003	0.012	0.007	0.012	0.017	0.006	0.014
	(0.009)	(0.015)	(0.022)	(0.015)	(0.012)	(0.005)	(0.011)
Observations	-0.000	0.483	0.000	0.471	0.224	0.036	0.185
Control mean	3334	3334	4206	4206	4205	4205	4205

Notes: **Table A35 shows that there are no average treatment effects on self-esteem or education investments in the big experiment.** Columns 1 and 2 show effects on self-esteem in the SMS survey 2-3 days after treatment on one item from the Rosenberg (1965) scale, "Sometimes I think I am no good". Jobseekers texted back one number "do you 4) strongly agree 5) agree a little 6) neutral 7) disagree a little 8) strongly disagree?". Columns 3 and 4 show effects on self-esteem measured in the endline survey 2-4 months after treatment on a five-item scale adapted from the Rosenberg (1965) scale, administered by fieldworkers as part of the phone survey. Each item measures, on a Likert scale, the extent to which jobseekers agree with a statement indicating high self-esteem. Columns 1 and 3 show effects on standardized measures and columns 2 and 4 show results for above median dummies. Column 5 shows the effect on a dummy indicating any newly started education. Column 6 shows the effect on a dummy indicating any newly started apprenticeship. Column 7 shows the effect on a dummy indicating any other newly started formal post-secondary education. Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A36: Heterogeneous Treatment Effects on Labor Market Outcomes by Skill Report Attachment to Applications - Big Experiment

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

Notes: **Table A36 shows that treatment effects on labor market outcomes in the big experiment are not driven by job-seekers that report attaching their skill reports to at least one application.** *Attached report w. application* is dropped from the regression as no control individual received a report. Columns 1 to 5 show impacts on employment quantity. Column 1 shows impacts on an Anderson (2008) index of the four employment quantity measures. Column 2 shows the impact on the number of job offers in the last 30 days (winsorized). Column 3 shows the impact on a dummy indicating any work for pay in month 1 after treatment. Column 4 shows the impact on a dummy indicating any work for pay in month 2 after treatment. Column 5 shows the impact on a dummy indicating any work for pay in the last seven days. Columns 6 to 9 show impacts on employment quality. Column 6 shows the impact on an Anderson (2008) index of the three employment quality measures. Column 7 shows the impact on earnings in the last seven days (winsorized). Column 8 shows the impact on hourly wages in the last seven days (winsorized). Column 9 shows the impact on a dummy indicating a written contract. Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). All monetary figures are reported in 2021 USD PPP. Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B Beliefs About Wages and Job Offer Probabilities

The model predicts a four-stage reaction to treatment: beliefs about skill comparative advantage $\rightarrow$ expected return to skill-directed search $\rightarrow$ search direction $\rightarrow$ search outcomes. We show evidence consistent with the first, third, and fourth stages in the paper. We assess the second stage in this appendix, using measures of jobseekers' expectations about their labor market outcomes.

We find treatment effects on beliefs about search outcomes are broadly consistent with the model. But we view them as less important than the skill belief, search direction, and search outcome results we include in the main paper. The questions about expected search outcomes rely on complex forecasts by jobseekers, as discussed in recent reviews (Delavande, 2022; Mueller and Spinnewijn, 2022). 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 rather than survey questions. And the survey questions about skill beliefs and current labor market outcomes are simpler to answer.

We only use data from the tight experiment in this section. In the tight experiment, we collect 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 occurs 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 we make in this paper.

Treatment effects on beliefs about returns to search direction: The model predicts that changes in search direction will be driven 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 collect two types of data.

First, we survey jobseekers about their expected outcomes from applying to each type of job and estimate treatment effects on these measures. We ask for the expected number of job offers in the next 30 days, time to employment, and wage when employed. We ask 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 ask all questions separately about all jobs, communication-heavy jobs, and numeracy-heavy jobs. We define the expected return to skill-directed job search as the expected outcome for jobs aligned to comparative advantage minus the expected outcome for jobs nonaligned to CA. For example, for a jobseeker with numeracy CA, the expected number of offers for numeracy jobs minus the expected number of offers for communication 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 (165% of the control group mean with $p = 0.023$) and this effect, like most in the paper, is driven by jobseekers with baseline misaligned CA beliefs (Table A37, columns 3-4). Treated jobseekers with baseline misaligned CA beliefs expect weekly earnings 24 USD higher from skill-directed job search: roughly four times the control group mean with $p < 0.001$ (column 8). Treatment has negligible effects on the expected length of time until getting a job from skill-directed job search (columns 5-6). We return to this result on page 89.

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 0.10 standard deviations ($p = 0.35$) for the average jobseeker (column 1). This result is driven by the same heterogeneity we see elsewhere in the paper: jobseekers with initially misaligned CA beliefs increase their expected returns by 0.26 standard deviations ($p = 0.054$) while jobseekers with initially aligned CA beliefs do not increase their expected returns (column 2).

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

We find similar but much less precisely estimated results using belief measures collected during the job choice task (Table A39). For 5 of the 11 job pairs, we ask jobseekers about the probability of getting an offer if they apply, 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. Within each pair of jobs, treatment increases the expected offer probability and wage for the job aligned with the jobseeker's comparative advantage (Table A40). But the effects are small – roughly 2% of the control group mean and 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 results using the survey measures of expected returns to skill-directed search discussed above because the questions 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, consistent with a role for belief-based job search. For 5 of the 11 pairs of jobs, we asked jobseekers for the probability 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 expected monthly wage using a logit regression model, following Wiswall and Zafar (2015). This is not an experimental analysis because we regress post-treatment choices on post-treatment beliefs. Column 1 of Table A38 shows that a 100 USD increase in the expected weekly wage scaled by the offer probability is associated with a 6.7 percentage point increase in the probability of choosing that job ($p = 0.009$). This relationship

is robust to adding both jobseeker and job pair fixed effects (columns 2-4). Treatment does not significantly change the slope of the belief-choice relationship, suggesting that at least part of the treatment effect on choices operates through beliefs.

Treatment effects on beliefs about search outcomes: The preceding analysis focuses on jobseekers' beliefs about the returns to skill-directed job search, which we define as the difference in jobseekers' expected outcomes from searching for jobs that do and do not align with their comparative advantage in skill assessments. 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 (5).

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 A40 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 7.4 and 14.9 respectively) perhaps reflecting the difficult 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).[35] 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 cannot 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 A41 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.[36] 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. But the magnitudes are substantially different, perhaps in part because the control

[35] In 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.

[36] For 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?

group's expectations are much higher than realized outcomes.

Table A37: Treatment Effects on Beliefs About Returns to Skill-Directed Job Search - Tight Experiment

	Index		ΔE [offers per app] (w)		- Δ [sear. dur.] (w)		ΔE [wage] (w)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	0.095	0.260*	0.043**	0.053**	-0.007	-0.079	0.073	23.981**
	(0.100)	(0.130)	(0.018)	(0.026)	(0.157)	(0.185)	(6.605)	(9.383)
Treatment × Aligned comp adv belief (bl)		-0.341*		-0.023		0.114		-45.738***
		(0.180)		(0.034)		(0.223)		(13.960)
Aligned comp adv belief (bl)		0.689***		0.107***		0.432**		35.845***
		(0.140)		(0.030)		(0.197)		(12.069)
Treatment effect: Aligned comp adv belief (bl)		-0.081		0.030		0.035		-21.757**
		(0.124)		(0.022)		(0.191)		(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 A37 shows that providing jobseekers with information about their comparative advantage in skills increases beliefs about the returns to directed search.** All outcomes are coded so that positive values mean higher beliefs about the returns to skill-directed search. Columns 1 and 2 show treatment effects on an index of belief about the returns to directed search. Columns 3 and 4 display effects on the difference in expected offers per application for jobs that align with jobseekers' comparative advantage and those that do not align with their comparative advantage. Columns 5 and 6 display effects on the difference in expected search duration for jobs between searching for jobs that *do not* align with jobseekers' assessed comparative advantage in and jobs that align with their assessed comparative advantage (winsorized). Columns 7 and 8 display effects on expected weekly wages for jobs that align with jobseekers' comparative advantage and those that do not align with their comparative advantage measured in 2021 PPP USD (winsorized). Even columns show heterogeneity by whether individuals have aligned comparative advantage beliefs at baseline. Controls include randomization block fixed effects, and prespecified baseline covariates (skill quintile dummies, age, gender, dummies for having completed a university degree, other post-secondary certificate or diploma, and high school, and a dummy indicating a comparative advantage in numeracy). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A38: Association Between Job Choices, Offer Probabilities, and Expected Wages - Tight Experiment

	Marginal effects on choice of numeracy job (logit estimation)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$E[Wage^{num}] - E[Wage^{com}]$ (w)	0.067***	0.075***	0.065**	0.075**	0.067***	0.077***	0.065**	0.073**
	(0.025)	(0.026)	(0.029)	(0.031)	(0.025)	(0.025)	(0.030)	(0.030)
Treatment × $E[Wage^{num}] - E[Wage^{com}]$ (w)					0.045	0.042	-0.017	-0.021
					(0.038)	(0.038)	(0.045)	(0.044)
# observations	925	925	800	800	1840	1840	1610	1610
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 A38 shows that, in the job choice task in the tight experiment, expected earnings are predictive of job choice and that treatment does not affect this relationship.** It displays the marginal effects of expected returns to applying for a job and job choices estimated using a logit estimator in the tight experiment. The outcome is a dummy indicating choosing the numeracy job in a job pair. The independent variable is the winsorized difference in expected returns to applying to the numeracy job and the communication job. Expected returns are defined as expected weekly wages (in 100s USD PPP in 2021) times the perceived likelihood of receiving a job offer. Columns 1 to 4 are limited to the control group. Columns 5 to 8 also interact the expected returns with a treatment indicator (columns 5 and 6 also include a treatment dummy alone). 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A39: 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	-3.240	-0.022
	(0.031)	(6.739)	(0.022)
Treatment× Aligned skill req.	0.008	3.873	0.011
	(0.030)	(3.491)	(0.015)
Aligned skill req.	0.030	-3.297	0.017
	(0.023)	(2.461)	(0.010)
Control mean	-0.000	193.588	0.544
Observations	2770	2770	2770

Notes: **Table A39 analyzes whether the treatment shifted beliefs about jobs differentially by whether skill requirements align with assessed comparative advantage.** Beliefs were elicited for 10 jobs in five job pairs for each jobseeker. Column 1 shows effects on a standardized measure of desirability of jobs measured on a 0 to 10 Likert-scale. Column 2 shows effects on expected weekly earnings winsorized at the 99^{th} percentile. Column 3 shows effects on the perceived job offer probability. Analysis is at the job-individual level and includes pre-specified control variables (skill quintile dummies, age, gender, and dummies for having completed a university degree, other post-secondary education, and high school) as well as randomization block, job choice order and job fixed effects. Standard errors clustered at the treatment-day level are in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A40: 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	0.319*	9.566	22.119	8.934	19.180*	0.650	3.115	1.318	3.805
	(0.104)	(0.179)	(7.740)	(14.852)	(6.090)	(10.607)	(11.984)	(24.259)	(3.628)	(5.948)
Treatment × Aligned comp adv belief (bl)		-0.461*		-26.041		-21.065		-6.425		-5.366
		(0.258)		(21.034)		(17.531)		(38.832)		(6.578)
Aligned comp adv belief (bl)		0.247		10.018		4.604		20.878		5.287
		(0.208)		(17.725)		(15.279)		(29.176)		(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 A40 shows that there are positive but insignificant average treatment effects on beliefs about wages in the tight experiment.** This effect is positive and significant for people with initially misaligned comparative advantage beliefs. Columns 1 and 2 show effects on an Anderson (2008) wage expectation index. All monetary figures are reported in 2021 USD PPP per week. (w) indicates measures winsorized at the 99^{th} percentile. Columns 3 to 10 show treatment effects on the individual index components. Columns 3 and 4 show impacts on expected wages. Columns 5 and 6 show impacts on the minimum expected wage. Columns 7 and 8 show impacts on the minimum expected wage. Columns 9 and 10 show impacts on the reservation wage. Even columns show heterogeneity by whether jobseekers had aligned comparative advantage beliefs at baseline. Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). Standard errors clustered at the treatment-day level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A41: 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	0.264	-0.001	0.039	-0.006	0.016	0.011	0.228	3.493	8.674**	4.235*	8.431**
	(0.122)	(0.194)	(0.023)	(0.036)	(0.021)	(0.031)	(0.247)	(0.392)	(2.238)	(3.748)	(2.315)	(4.042)
Treatment × Aligned comp adv belief (bl)		-0.360		-0.086		-0.049		-0.444		-10.821*		-8.678*
		(0.247)		(0.057)		(0.042)		(0.470)		(6.003)		(5.126)
Aligned comp adv belief (bl)		0.234		0.061		0.049		0.069		4.789		2.647
		(0.228)		(0.043)		(0.036)		(0.386)		(4.416)		(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 A41 shows that there are positive but mostly insignificant treatment effects on beliefs about job offers and search duration in the tight experiment.** Columns 1 and 2 show effects on an Anderson (2008) labor market expectation index. Columns 3 to 12 show effects on the individual index components. Columns 3 and 4 show impacts on expected callbacks per application. Columns 5 and 6 show impacts on expected offers per application. Columns 7 and 8 show impacts on the expected time until jobseekers find a full-time job (this enters the index negatively). Columns 9 and 10 show impacts on the perceived probability of being employed one month after baseline. Columns 11 and 12 show impacts on the perceived probability of being employed three months after baseline. (w) indicates measures winsorized at the 99^{th} percentile. Even columns show heterogeneity by whether jobseekers had aligned comparative advantage beliefs at baseline. Controls include randomization block fixed effects, and prespecified baseline covariates (measured skills, self-reported skills, education, age, gender, employment, discount rate, and risk aversion). Standard errors clustered at the treatment-day level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C Willingness-to-Pay Appendix

As part of the job search workshop we jobseekers' willingness-to-pay for three products: a document that reveals the expert-assessed skill requirements of jobs included in the job choice task (delivered after 30 days to avoid it influencing their search behavior), printed self-study materials for to improve numeracy skills, and printed self-study materials for to improve written communication skills. The self-study materials were taken from the "Mind the Gap!" series which targets grade 12 students and was developed by the Department of Basic Education in South Africa.

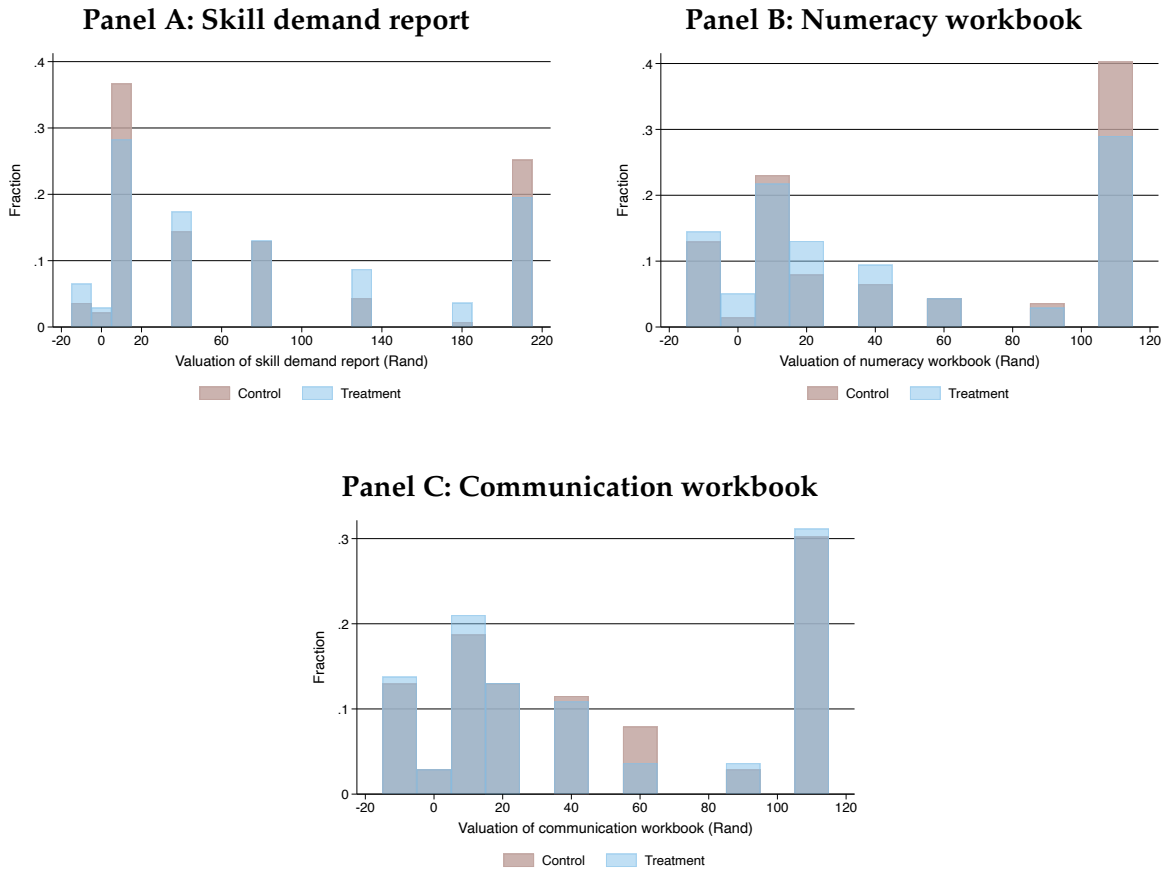
We elicit willingness-to-pay using multiple price lists with six items and announce that one randomly selected choice will be implemented in practice. The choice is always between receiving the product and receiving a changing monetary amount in airtime for their mobile phone (the "price"). For the skill document revealing the skill requirements of jobs the prices are: 0, 25, 50, 100, 150, and 200 Rand. For the self-study materials, we use the following prices: 0, 15, 30, 50, 75, and 100 Rand. We randomize whether jobseekers are asked about the highest or lowest price first. We force a unique switching point by ending the elicitation for an item at the first question they choose the item (for those starting with the highest price) or the monetary amount (for those starting with the lowest price) and assume that they would choose consistently on the omitted questions. We then randomly choose one question among the 18 questions and implement their choice in practice and announce this at the end of the session.

To ensure that jobseekers are familiar with the procedure, they complete a practice round where we elicit the jobseekers' willingness-to-pay for a bar of soap. If their choices are inconsistent, the enumerator explains the elicitation again and asks for the reason for the inconsistency to ensure that they understand the process.

To construct monetary willingness-to-pay measures, we assign jobseekers the average of the price they switched and the previous price. For example, if a jobseeker prefers the item at price zero but the money at price 25 Rand we assign them a willingness-to-pay of 12.5 Rand. For jobseekers who switch immediately, we assign -10 Rand (for those with negative willingness-to-pay) and 110 or 210 Rand (for those whose willingness-to-pay exceeds the maximum price). Given that there is substantially more mass on the upper end of the willingness-to-pay range, this conservative assumption is likely to lead us to underestimate the true willingness to pay. Figure A5 displays the distribution of willingness-to-pay measures by treatment status. Our treatment results are qualitatively robust to using ordered logit estimation rather than assigning monetary values.

Table A42 shows treatment effects on the three different willingness-to-pay measures. The average willingness-to-pay for the skill report in the control group is 79 Rand (11.06 USD). It is 55 Rand (7.7 USD) for the numeracy materials and 48 Rand (6.72 USD) for the communication materials. We observe no treatment effect on jobseekers' willingness-to-pay for the skill demand report. We find negative treatment effects on jobseekers' willingness-to-pay for the numeracy workbook, driven by individuals with a comparative advantage in numeracy and initially misaligned comparative advantage beliefs. We find negative treatment effects on jobseekers' willingness-to-pay

Figure A5: Distribution of Willingness-to-pay Measures - Tight Experiment



Notes: **Figure A5 shows the distribution of elicited willingness-to-pay for three goods.** Panel A displays the distribution of willingness-to-pay for a report revealing the skill demand of jobs in the jobs choice task. Panel B displays the distribution of willingness-to-pay for a numeracy workbook. Panel C displays the distribution of willingness-to-pay for a communication workbook. All values are in 2022 South African Rands to illustrate the decisions as made by jobseekers. All other currency values in the paper are reported in 2021 USD purchasing power parity at the exchange rate 1 Rand = 0.140 USD PPP.

for the numeracy workbook, driven by individuals with a comparative advantage in numeracy and initially misaligned comparative advantage beliefs and jobseekers with a comparative advantage in communication and initially aligned comparative advantage beliefs. We find no average treatment effects on jobseekers' willingness-to-pay for the communication workbook with some indication that jobseekers with a comparative advantage in numeracy and initially aligned comparative advantage beliefs increase their willingness-to-pay. Note that the cell sizes for heterogeneity analysis are relatively small which limits the inference one should draw from these analyses.

Table A42: Treatment Effects on Willingness-to-pay - Tight Experiment

	Info on skill requirements						Numeracy materials						Communication materials					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
	Pooled		Num. CA		Comm. CA		Pooled		Num. CA		Comm. CA		Pooled		Num. CA		Comm. CA	
Treatment	-0.719	-6.851	-12.823	-14.035	1.708	-8.179	-13.410***	-15.309**	-14.527*	-20.583**	-10.046**	-3.217	-0.822	-0.247	5.252	-2.571	-3.821	3.862
	(5.766)	(11.255)	(13.609)	(15.888)	(8.396)	(17.726)	(3.539)	(6.775)	(8.049)	(8.944)	(3.898)	(11.201)	(3.790)	(6.747)	(9.785)	(12.162)	(5.882)	(11.719)
Treatment × Aligned comp adv belief (bl)		12.287		8.804		14.077		5.019		36.732**		-9.825		-0.847		45.821**		-11.987
		(18.773)		(38.367)		(22.913)		(11.915)		(15.806)		(16.400)		(10.474)		(21.394)		(14.833)
Aligned comp adv belief (bl)		0.174		-12.737		7.989		-14.744		-38.444***		-4.144		-3.753		-43.991***		7.831
		(16.202)		(28.007)		(16.687)		(8.948)		(11.399)		(11.074)		(9.163)		(11.637)		(10.853)
Treatment effect:		5.436		-5.231		5.898		-10.290		16.149		-13.042*		-1.094		43.251**		-8.125
Aligned comp adv belief (bl)		(10.847)		(33.252)		(11.314)		(7.268)		(14.781)		(7.247)		(6.130)		(16.512)		(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 A42 analyzes how treated jobseekers in the tight experiment change their willingness-to-pay for different goods.** Columns 1 to 6 show effects on the willingness-to-pay for a document with the expert-assessed skill requirements for the 11 job pairs in the job choice task. Columns 7 to 12 show impacts on willingness to pay for numeracy course materials. Columns 13 to 18 show impacts on communication course materials. Columns 1, 2, 7, 8, 13, and 14 show results for the full sample. Columns 3, 4, 9, 10, 15, and 16 show results for individuals with a comparative advantage in numeracy. Columns 5, 6, 11, 12, 17, and 18 show results for individuals with a comparative advantage in communication. Controls include randomization block fixed effects, and prespecified baseline covariates (skill quintile dummies, age, gender, and dummies for having completed a university degree, other post-secondary education, and high school). Standard errors clustered at the treatment-day level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

D Skill Appendix

D.1 Big Experiment

In the big experiment we measured six skills: communication, numeracy, concept formation, focus, grit, and planning. Jobseekers completed the assessments at the Harambee office on a computer. Two or three industrial psychologists supervised the assessments and provided a debrief after the assessments were done.

A subset of the assessments (communication, concept formation, and numeracy) has been used by our partner to screen jobseekers in the past. Our partner had been contracted to screen roughly 160,000 prospective workers using these assessments by 2016 by firms in South Africa. Based on this, we presume that the information content of the assessments is valuable for the firms. This, however, does not mean that assessment results are the only information firms use in their hiring decisions. Additionally, 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 for jobseekers' productivity.

The *communication* test quantifies jobseekers' listening and reading comprehension. The assessment was developed by a local adult education provider (www.mediaworks.co.za) and focuses on high school English proficiency.

The *numeracy* assessment measures practical arithmetic skills and pattern recognition. The first part of the 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 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 very similar test to Raven's matrices (Raven and Raven, 2003). With this measure, we proxy for jobseekers' fluid intelligence (people's conceptual reasoning and the rate at which people learn).

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 and DiGirolamo, 1998).

We rely on Duckworth's self-reported 8-item scale to assess jobseekers' level of *grit* (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 point to a point basket. The party wins, whose action leads to a score of 15 in the point basket.

Appendix Table A43 shows the correlation matrix between different skill terciles in the field experiment. Scores are weakly correlated across assessments, with pairwise correlations between 0.09 and 0.44. Hence, the assessments horizontally differentiate candidates based on their relative

skills rather than only ranking or vertically differentiating them in a single dimension of skills.

Table A43: Correlation Between Skill Terciles - Big Experiment

	Concept formation	Communication	Numeracy	Grit	Planning	Focus
	(1)	(2)	(3)	(4)	(5)	(6)
Concept formation	1.000	0.298	0.435	0.156	0.120	0.140
Communication		1.000	0.331	0.111	0.094	0.119
Numeracy			1.000	0.172	0.124	0.116
Grit				1.000	0.157	0.168
Planning					1.000	0.132
Focus						1.000

Notes: **Table A43 shows positive but moderate Pearson correlation coefficients between the skill terciles used in the field experiment.**

D.2 Tight Experiment

In the tight experiment, we assess jobseekers on numeracy, communication, and general intelligence. The numeracy test is the subset of questions from the field experiment aimed at testing suitability for retail jobs. The communication test is equivalent to the test used in the field experiment. The test of general intelligence is based on 20 Raven's matrices. The assessment was completed in the field location using tablets.

Table A44: Correlation Between Skill Quintiles - Tight Experiment

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

Notes: **Table A44 shows positive but moderate Pearson correlation coefficients between the skill quintiles in the tight experiment.** Panel A shows results for the sample with a clear comparative advantage between numeracy and communication. Panel B shows results for the full sample.

Table A44 shows the correlations between skill measures in the tight experiment. We find moderate positive correlations between all three skill measures. Table A45 displays correlations

between pre-treatment beliefs about skills and assessment results in the experiment. We find that numeracy beliefs are strongly correlated with assessment results but communication skills are not significantly correlated with assessment results.

Table A45: Association Between Assessed and Believed Skill Quintiles - Tight Experiment

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

Notes: **Table A45 analyzes the correlations between assessed skill quintiles and baseline skill quintile beliefs in the tight experiment.** No control variables are included. Robust standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

E Benefit-Cost Comparison

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

The average treatment effect on weekly earnings is 6.52 USD at the time of the follow-up survey, which occurs 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 on assessment and software licenses; 8.11 on airtime and data, mainly for internet access for the assessment computers and for contacting jobseekers; 3.82 on Harambee salaries for assessment support staff, psychologists delivering results briefings, and administrative support; 1.40 on J-PAL Africa salaries for research staff who helped to run the assessments and results briefings; and 9.74 on transport money for jobseekers to attend the assessments.

This implies that the ratio of average benefit to average variable cost is 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 like facility and equipment rental, most staff costs, and assessment and software licenses. The only costs 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 but automated feedback on assessment results. These involve large fixed costs of development but they reduce or eliminate most of the large components of the average variable of in-person assessment: facility rental (27% of average variable cost), equipment rental (22%), and transport (21%).

F Preregistration Appendix

The paper relies on two pre-registrations and analysis plans (AEARCTR-0001631 for the big experiment and AEARCTR-0010000 for the tight experiment).

The pre-analysis plan for the big experiment includes the experimental design, the treatment effect analysis, and the main labor market outcomes. As the exploratory analysis of the big experiment lead us to the research question of the current paper, the pre-analysis plan did not include the specific research question of the current paper. We deviate from the plan in the following way:

- We restrict the sample to the *private* treatment arm. Participants in this arm received the unbranded report without any identifiers about their relative performance in terciles.
- We study additional outcomes on the topic of comparative advantage beliefs and aligned search. These outcomes were not prespecified.
- The treatment effect analysis was modified to reflect the change in the sample and to study heterogeneous effects by aligned comparative advantage.

The tight experiment was set up to test the research question that arose from the exploratory analysis and thus the pre-analysis plan is more exhaustive. We deviate from the pre-analysis plan in the following way:

- We restrict the sample to people who have a clear comparative advantage. In Appendix Tables A8 and A9 we show the main results of the paper for the full sample as a robustness check. The interpretation of results remains unchanged.
- We added further outcomes on search alignment in the tight experiment. These were SMS and platform outcomes (and their index) that we were unsure about whether we can get access to from our partner at the time of the pre-registration. We correct for the inclusion of these additional measures by creating a summary search alignment index.
- To align the search alignment 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 measure. We show results for the pre-specified platform search index in Table A32.
- Following recent methodological critiques (Chen and Roth, 2022; Mullahy and Norton, 2022) we use winsorization instead of the pre-specified inverse hyperbolic sine transformation to handle outliers. The key results are all robust to using the inverse hyperbolic sine instead of winsorization.

G Protocol and Intervention Details

G.1 Tight experiment

The data collection for the tight experiment ran between August and October in 2022 in downtown Johannesburg. We recruited 373 participants using the user database of our implementation partner. We contacted users who were active on SAYouth.mobi in the past month, said they were residing in Gauteng province, completed at least a high school leaving exam and were weakly less than 35 years old. Using this contact list, surveyors called potential participants, and after a short screening[37], invited them for a daylong job search workshop in the city center. We offered 150 Rand (approximately 21 USD) airtime for compensation.

When participants arrived at the venue, participants 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 to the participants 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, and also measured risk and time preferences. This survey was followed by three assessments; participants completed the communication, numeracy, and cognitive performance tests in order. Participants had 30 minutes for each of the communication assessment and numeracy assessment and 15 minutes to complete the cognitive performance test. Section D describes the skills and tests in more detail. 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 any questions from the surveyor. On control days, participants still viewed a minimally modified version of the video that omitted the explanation of the results. The scripts and the video was thoroughly piloted to ease participants' understanding.The treatment video is available at https://www.dropbox.com/s/9qug1673wxefdxj/treatment_video_final.mp4?dl=0 and the control video is available at https://www.dropbox.com/s/t4d57xmwopy8r4q/control_video_final.mp4?dl=0.

After the treatment and a lunch break, the surveyors elicited 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. The job titles 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. The job titles were all entry-level jobs that did not require certifications or

[37]The screening questions confirmed if the user is currently looking for a job, not a full-time student, and has not visited the lab in-the-filed before (even during the piloting).

equipment to ensure that all participants could reasonably apply for them (see Table 5). The job descriptions and their layout followed the design and style of the SAYouth.mobi platform. The job descriptions were presented in a printed booklet side-by-sideto allow for easy comparisons (see Figure 2). Participants were shown the pairs in the booklet, read the descriptions and were asked to pick the job that they are 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 (numeracy or communication) that the job required. When participants completed their choices, we elicited their beliefs and experience related to each job in 5 out of 11 pairs.

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 it is during the workshop. Second, participants received a list of job titles 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 for three products described in section C. We then collected information typically found in CVs. We used this information to create a CV for each participant. We prepared and delivered the CVs to participants after the workshop, and we 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 job in the job choice task that they preferred. The message, along with participants' CV that we created, were both delivered to the respective employers. After this part, participants completed the check out procedure and received their compensation.

G.2 Big Experiment

The recruitment, data collection procedure and intervention of the big experiment is very similar of the tight experiment. The specific details are described in Carranza et al. (2022), in this section we highlight the main differences between the big and tight experiment.

The big experiment took place in the same labor market, but earlier in time, in 2016/17. We again used the contact list of the Harambee Youth Employment Accelerator. Since the job search platform was not yet developed, we recruited young workseekers (aged 18–29) who have completed secondary school, have at most 12 months of formal work experience and are from disadvantaged backgrounds. The venue of the data collection was the Harambee office at the time in downtown Johannesburg.

Similar to the tight experiment, workseekers completed surveys and assessments. Treatment was randomized at the day level and the intervention consisted of the receiving a report. The report (Figure A3) 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 a psychologist. The content of the briefing is akin to that in the tight experiment video. The materials used for the briefing are available at https://www.dropbox.com/s/b3lumodiz1gxadq/4c_Presentation_Private_24Oct2016.pptx?dl=0 and the script is available at https://www.dropbox.com/s/4lhxf502sy5kw3r/4d_Private%20Briefing_24Oct2016.docx?dl=0.

Two follow-up surveys were conducted. A short SMS survey 1-3 days after the workshop, and a longer phone survey between 3-4 months after the workshop.

H Firm Preferences Over Skills

We use an incentivized choice experiment to show that firms' vary in their valuation of communication and numeracy and value both highly relative to some forms of education. Carranza et al. (2022) describe the sample and data collection in detail and use this to show that there is substantial variation in firms' preferences over skills.

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. They 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.

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 profiles with different levels of skills and with or without a post-secondary diploma, all with completed secondary school. To incentivize the choices, firms' rankings were used to match them with jobseekers with specific skill profiles from Harambee's database, in a similar spirit to Kessler et al. (2019).

Table A46 shows the ranking of numeracy, communication and education, averaging over the 67 firms. There are six different possible rankings of these three elements, each shown in a row. Column 6 shows that 57% of firms prefer a candidate with top-tercile numeracy skills, 34% prefer a candidate with top-tercile communication skills and 9% firms prefer a candidate with middle-tercile communication and numeracy skills and with a diploma.

Table A46: 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	Com	Educ	52.24%	Numeracy	56.72%
2	Num	Educ	Com	4.48%		
3	Com	Num	Educ	28.36%	Communication	34.33%
4	Com	Educ	Num	5.97%		
5	Educ	Num	Com	1.49%	Educaction	8.95%
6	Educ	Com	Num	7.46%		
				100.00%		100.00%

Notes: **Table A46 shows that firms, on average, value applicants' numeracy (Num) and communication skills (Com) 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. Column 4 shows the share of firms who chose the respective rankings. Column 5 and 6 does the same collapsing rows according to the most important skill.