Essays on Labor Markets in Developing Countries

Essays on Labor Markets in Developing Countries

A dissertation presented

by

Simon Martin Abel

to

Graduate School of Arts and Science

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Public Policy

Harvard University

Cambridge, Massachusetts

April 2017

Dissertation Advisors:
Professor Rema Hanna
Professor Lawrence Katz

Author:
Simon Martin Abel

Essays on Labor Markets in Developing Countries

Abstract

Labor market frictions are particularly prevalent in developing countries. My dissertation documents the extent of various market frictions, investigates its economic consequences, and tests interventions aimed at addressing these frictions in the context of South Africa. Chapter 1 documents substantial information asymmetries about workers' productivity between job seekers and hiring firms. We design and experimentally test a reference letter and find that it leads to employment gains and an overall increase in match quality. Chapter 2 investigates the problem that job seekers fail to follow through on their job search plans, which reduces the available applicant pool for firms. We conduct a field experiment and find that a simple planning intervention implemented as part of at government-run job counseling workshop leads to a significant increase in job search and job opportunities. Chapter 3 explores discrimination among hiring firms. I collect survey data and a unique data set of classified ads and employ a novel quasi-experimental method that exploits variation in the applicant pool composition. Results show that firms discriminate against immigrant workers and that these job seekers respond by adjusting their search strategy.

Contents

List of Tables

List of Figures

Acknowledgements

I am extremely grateful to my advisers Rema Hanna, Larry Katz, Asim Khwaja and Michael Kremer for their continuous support, guidance and encouragement. My work also benefited from many discussions with Simeon Abel, Dany Bahar, Willa Brown, Rulof Burger, Eliana Carranza, Shawn Cole, Jasmina Chauvin, Juan Pablo Chauvin, Ben Duclos, Raissa Fabregas, Leander Heldring, Tomoko Harigaya, Ed Kerby, Sara Lowes, Samantha Ndiwalana, Oyebola Okunogbe, Patrizio Piraino, Philipp Werner and Bilal Zia. I am also thankful for the cooperation with the South African Department of Labour and their support in the implemention of my research projects.

To Willa B.

In Memory of Thomas

Introduction

About 80% of the global work force lives in the developing world. Gary Fields notes that "*the main and often the sole asset of the poor is their labor. So to understand global poverty, one must understand labor markets and labor earnings in the developing world*" (Fields (2011). My research focuses on the matching process between workers and firms using the case of South Africa, a country with one of the highest unemployment rates in the world.

It is widely believed that frictions in the matching process are more prevalent in developing countries. For example, information asymmetries may be particularly large as there are fewer institutions who can credibly certify the skills of workers. Education systems are often of lower quality which limits the reliability of educational degrees as productivity signals. Chapter 1 of my dissertation tests an intervention to reduce these information asymmetries. A second friction stems from inefficient job search among the unemployed. A failure to submit (informative) applications reduces the applicant pool available to hiring firms, which may result in lower match quality. Chapter 2 explores an action plan intervention that tries to address behavioral barriers to effective job search. Thirdly, discriminatory hiring practices undermines an allocation of jobs based on productivity. This may be particularly prevelant in the informal sector, where vulnerable groups are less protected by anti-discrimination regulation. Chapter 3 investigates discrimination towards immigrant workers in the context of South Africa's informal sector.

The overarching goal of this research is to use experimental and quasi-experimental methods to provide policy-relevant evidence. Two of my chapters resulted from research projects conducted in close cooperation with the South African government. The interventions I tested - reference letters and a plan making prompt for job seekers - will be scaled up and integrated into the government's employment services. Next, I provide a brief summary of the main results of each chapter.

Chapter 1 explores whether reference letters from former employers, which are rarely used in developing country settings, could play a role in alleviating information asymmetries. We conduct a series of field experiments to investigate the value and usage of standardized reference letters among young job seekers in South Africa. A resume correspondence study finds that including a reference letter with the application increases employer call-backs by 60%. Women, traditionally excluded from many referral networks, particularly benefit: firms pay closer attention to the content of letters sent by women and increase response rates by 89%. A second experiment, which encourages job seekers to obtain and use a reference letter, finds similar results. Men are not more likely to find jobs, but employment rates for women who have reference letters double, thus fully closing the employment gender gap in our sample after three months. We find that letters are effective because they provide accurate information about workers' skills that firms use to select applicants of higher ability, unless they deem letters to be implausibly positive. Despite these positive findings, reference letters are not widely adopted, partly because job seekers underestimate their potential value.

Substantial research in behavioral economics and psychology documents an "intention-behavior gap": an imperfect relationship between intentions and follow-through. Chapter 2 extends this research to the novel domain of job search. We present a field experiment with unemployed youths in South Africa to test the effect of an action planning intervention on job search and employment. Five to twelve weeks after the intervention, the action planning treatment increases the number of job applications submitted (15%), but does not affect total job search hours. This is consistent with the intention-behavior gaps measured at baseline and suggests improved efficiency of search. In addition we find that job seekers diversify their search strategy and use more formal search channels. These increases in search efficiency and effectiveness translate into more job offers (30%) and employment (26%). We further combine action planning with weekly reminders and a peer-support intervention. Neither of these additional interventions improves the effectiveness of action planning, providing suggestive evidence against commitment and limited attention as underlying mechanisms.

In Chapter 3, I use a unique data set of classified ads in South Africa to explore whether employers discriminate against immigrants in the hiring process. I develop a quasi-experimental method to estimate discrimination exploiting variation in the applicant pool composition due to the timing of postings. Consistent with a tournament models in which immigrants are penalized, I find that both foreigners and natives benefit from being pooled with foreign job seekers. Next, I test whether discrimination affects search behavior. Controlling for location fixed effects, I find suggestive evidence for sorting: immigrants search further away and higher discrimination in the residential area is positively correlated with the decision to search in different suburbs. This additional cost to job seekers has not been explored in the discrimination literature.

1 The Value of Reference Letters - Experimental Evidence from South Africa

1.1 Introduction

[1]Information asymmetries about workers' skills are prevalent in labor markets, especially in the market for low skill and entry level jobs.[2] In developed economies, hiring firms commonly reduce these asymmetries through reference letters from previous employers (Ioannides and Loury, 2004). In developing countries, this practice is largely absent and many firms instead hire through informal referrals, such as those from their existing work force.[3] This has potential adverse effects on match quality as it limits the pool of candidates (Loury, 2006) and as current employees may refer close friends or family members rather than their most qualified peers (Beaman et al., 2013). In addition, informal referral systems may exacerbate inequity as they disadvantage less connected groups; in particular, they harm women who often lack access to informal referral networks (Beaman et al., 2013; Montgomery, 1991).[4]

There are various potential reasons why reference letters are not more widely used in many markets: former employers may not be willing to provide relevant information, hiring firms

[1]This paper is co-authored with Rulof Burger (Stellenbosch University) and Patrizio Piraino (University of Cape Town).

[2]In these markets, job seekers often have limited work experience and lack educational degrees to signal skills. Firms are less likely to invest in costly screening as employment relationships are often short-term (Autor and Scarborough, 2008). A literature on firm learning provides indirect evidence that information asymmetries are prevalent at the time of hiring: returns to easily observable characteristics (e.g. race, education) diminish as employers learn about initially unobserved characteristics over time (Kahn and Lange, 2014; Farber and Gibbons, 1996; Altonji and Pierret, 2001).

[3]In developed economies about 50% of jobs are found through informal network (Topa, 2011) compared to about 68% in South Africa (Schoer et al., 2014).

[4]In South Africa, many young job-seekers do not have strong connections to the labor market via employed friends and family members. Female job seekers may be at a particular disadvantage; previous research shows that women are more reliant on social networks and informal channels in the search process (Schoer and Leibbrandt, 2006) and that family networks in South Africa favor male members (Magruder, 2010).

may not perceive letters to be credible or informative, and job seekers may not request them as they underestimate their value or procrastinate.

We conduct three experiments in cooperation with the South African Department of Labour (DoL) to distinguish between these explanations. Specifically, we design a reference letter template and encourage young job seekers to have a former employer complete it. To test whether reference letters are valuable in principle, we first submit applications on behalf of job seekers to vacancies with and without reference letters and compare firm responses (Exp.1).[5] To assess whether letters are valuable in practice, we conduct an additional experiment in which we encourage half of job seekers to obtain a letter and subsequently follow their job search behavior and employment outcomes (Exp.2). Given the large positive effects we find, we run a third experiment that tests different explanations for why letters are not more prevalent (Exp.3). Evidence from these experiments enables us to answer three questions: i) Do reference letters have value? ii) How do they generate value? and iii) What explains their (lack of) usage?

We find that reference letters are valuable to both job seekers and hiring firms. Attaching a letter increases the probability that a firm responds to the applicant by 59% (from 4.2% to 6.5%) and the rate of interview requests by 60% (from 2.2% to 3.6%). Effect sizes are larger for women, increasing employer responses by 89% (results from Exp.1). While we do not detect significant impacts for men, female participants who obtained letters are more likely to receive job interviews and their employment rate doubles after three months, thus fully closing the gender employment gap in our sample (Exp.2). On the firm side, reference letters help to select candidates of higher ability: performing one standard deviation higher on an aptitude test increases the likelihood of an employer response by 0.6 percentage points (15%) for applications that do not include a reference letter (Exp.1). Attaching a letter increases

[5]We are among the first to conduct an audit study with actual job seekers. This addresses the criticism that application materials designed by researchers may not be realistic or include all relevant information (Heckman, 1998) as well as ethical concerns about sending fictitious applications (Riach and Rich, 2004).

this figure to 2.6 percentage points (63%).[6] These improvements in firms' screening ability apply to both male and female candidates.

How do reference letters generate value? We find that letters are informative of workers' skills: ratings from previous employers are highly correlated with aptitude scores of both male and female job seekers, even after controlling for information that can be easily inferred from the resume or school transcripts. Firms correctly use this information to update their beliefs of applicants and are more likely to respond to applications with positive letters (Exp.1). However, reference letters in which the former employer gives the highest rating in *every* category are completely ineffective, despite the fact that job seekers with these glowing reference letters perform very well in the aptitude test. This suggests that a perceived lack of credibility of the letter harms their employment prospects. The effect of employer ratings is more pronounced for women, indicating that firms are more uncertain about their skills and thus pay more attention to their letters' content (Exp.1).

In light of these results, one might ask why reference letters are not more prevalent and those that do exist often lack relevant information.[7] At baseline, 88% of job seekers say they do not have a letter because they "never asked", often claiming they did not know they needed one. Once encouraged, 31-42% of participants succeed in obtaining a letter. This share increases substantially when we provide job seekers with information on the benefits of having a letter (Exp.3). By contrast, an arm of the experiment which offered cash payouts for obtaining letters had no effect. Underestimating potential benefits may thus explain why job seekers are not asking former employers to provide (informative) reference letters. To test why job seekers do not discover their effectiveness, we analyze actual applications submitted

[6]Although our study was not designed to explicitly test for general equilibrium effects, theory predicts that these reductions in information asymmetries increase firm demand (Pissarides, 1985).

[7]Less than 2% of people in the control group use a reference letter in the job search and the majority of these letters are generic and do not provide information about skills (72%) or even the reference's contact information (44%). Interviews with firms indicate that employers know what content a letter should include, but providing this information to the market is costly and references do not directly benefit from it (Avery et al., 1999).

by participants and find that only about 20% of people who obtain a letter submit it as part of their application (Exp.2). This low usage stifles any learning about their benefits through job search. We also find that among participants in the treatment group, women are significantly more likely to use letters in their search, which explains part of the large gender difference in employment effects.

Given the low prevalence of reference letters in the South African labor market for low-skill jobs, we explore how their effectiveness may differ as letters become more common. First, we introduce a simple model of employer learning that predicts that the penalty of *not* sending a letter increases as reference letters become more prevalent. Simulation results confirm that this in turn induces more job seekers to submit a letter, improving the ability of firms to screen applicants. A second reason that the effect may differ if reference letters are more common is that an application is less distinguished by having a letter attached. We test this hypothesis by randomly varying the share of applicants for whom we submit reference letters. We find that increasing the number of reference letters sent to a given vacancy from one to three does not affect the letters' impact (Exp.1).

This study contributes to the literature on job referrals. Previous studies have largely focused on whether social network links can be exploited to reduce information asymmetries, showing that although workers have information on the productivity of their peers (Pallais and Sands, 2016; Burks et al., 2015), they are less likely to pass on truthful information to firms unless sufficiently incentivized (Beaman and Magruder, 2012). Former employers may provide more credible information because their incentives are more aligned with the hiring firm and they can assess workers more accurately as they observed them in a professional setting (Aamodt, 2015). However, few studies have looked at the role of former employers in reducing information asymmetries. Two notable exceptions are Pallais, 2014, who finds that feedback on workers' past performance in an online labor market increases their employment prospects, and Bassi and Nansamba, 2017 who study the effect of certifying soft skills. We contribute

to this literature by investigating a more traditional labor market setting in which workers can choose both the referee and whether to reveal the information to the market after they observe it.

In addition, we contribute to the literature on how search frictions affect employment (Mortensen and Pissarides, 1994). Information asymmetries between firms and workers lead to socially sub-optimal hiring of people without work experience and an overall decrease in market efficiency (Pallais, 2014; Terviö, 2009). We identify an additional market inefficiency: *after* people are hired, firms lack the incentive to provide (detailed) information about the worker's ability to the market.[8]

This study also adds to an extensive literature evaluating the effectiveness of active labor market policies (ALMP) (see Card et al., 2015 and McKenzie, 2017 for recent reviews). The evidence on ALMPs is mixed, in part because they typically include a package of interventions which makes it difficult to isolate the effectiveness of specific components. In this study, we are able to isolate one component of ALMPs, namely the reduction of information asymmetries.[9]

Lastly, we are to our knowledge the first to experimentally test the effect of reference letters, a ubiquitous selection tool in many developed economies, on employment.[10] Our results

[8]This provides a rationale for the government to intervene and facilitate this information exchange. The South African government introduced a wage subsidy for unemployed youths in 2014 which may address the inefficiently low hiring of job seekers without work experience. The lack of information exchange between former employers and hiring firms reduces the effectiveness of this policy, especially given that many of the employment opportunities are short term.

[9]Two recent studies test the effect of reducing information frictions by a third party in the context of developing countries. In Jordan, Groh et al., 2014 use results from psychometric and skill testing to match job seekers to vacancies. Test results are predictive of subsequent job market outcomes, but the intervention did not lead to changes in employment as 83% of job seekers rejected job offers or quickly quit. Abebe et al., 2016 test a combination of job counseling and skill certification for job seekers in Ethiopia. While overall employment and wages are unchanged, they find a large increase in permanent employment.

[10]There is relatively little research on reference letters, defined as a "description or evaluation of an applicant that is completed by an observer and used as a source of information for personnel selection" (McCarthy and Goffin, 2001), despite its ubiquity in the selection process (Aamodt, 2015). Existing research focuses on the ability of reference letters to predict future performance. One exception is Kaas and Manger,

suggest that letters can also be effective in developing country contexts: they both enhance firms' screening ability and benefit job seekers. In particular, we find large employment gains for women, a group often excluded from informal referral networks. Reducing information asymmetries - through reference letters or other interventions - may thus improve equity by leveling the playing field for women in labor markets.

The remainder of this paper is structured as follows: Section 1.2 describes the study context and introduces a conceptual framework. Section 1.3 describes the study designs. Section 1.4 investigates the effects of reference letters on firm responses and employment. Section 1.5 explores how firms use reference letters in the hiring process. Section 1.6 discusses why reference letters are not more widely adopted and Section 1.7 concludes.

1.2 Background and Conceptual Framework

1.2.1 South Africa's Labor Market

The post-2008 economic slowdown coupled with the rapidly growing working-age population contributed to a persistently high unemployment rate in South Africa (26.4%), especially for youths (36.9%) (StatsSA, 2015). The gender gap among black South Africans is substantial, despite the fact that black females are on average more educated than their male counterparts (Rospabe, 2001; Shepherd, 2008). One explanation is that firms appear to either underestimate or are more uncertain of the ability of female applicants.[11]

2012 who find through an audit study that reference letters do not increase overall employer responses but may benefit applicants from minory groups. Most studies find that reference letters are only moderately predictive of performance, in part because employer assessments are overly positive (McCarthy and Goffin, 2001).

[11] Malindi, 2016 finds that black females have a much higher returns to job tenure than black males, white females or white males in South Africa. This is consistent with a model in which employers initially underestimate or attach greater uncertainty to the value of productive attributes possessed by black females, but then upwardly adjust their wages once they observe their true productivity.

South Africa offers a context conducive to investigating the role of labor market information asymmetries. Most of the unemployed did not complete secondary education (55%) and have no or limited work experience (50.6%), which leaves firms with very little information to screen job applicants. In addition, the quality of education is perceived to be low in the majority of South African schools which limits the use of educational credential as signals for productivity (van der Berg, 2007). Last, in economies with mass unemployment, the employment status is less indicative of job seekers' ability (Kroft et al., 2013).

Information asymmetries affect how firms and workers are matched. In South Africa, some large firms administer aptitude tests as part of the hiring process. While these tests can increase aggregate productivity and labor demand by improving match quality (Mortensen and Pissarides, 1994; Pissarides, 1985), they have not been widely adopted for several reasons: First, when faced with hundreds of applications firms still need to first decide who to test and may therefore overlook the most suitable candidates.[12] Second, Autor and Scarborough, 2008 note that firms have fewer incentives to test candidates for jobs where investment in training is limited and employment spells are brief. Last, many small firms lack the expertise and resources to systematically test applicants.

Faced with these challenges, South African employers have increasingly turned to social networks and the existing workforce to fill vacancies.[13] Yet, firms face a trade-off in their choice of hiring channels (Montgomery, 1991). Under the "good match" hypothesis (Rees, 1966), current employers can help overcome the asymmetric information problem and create better employment matches as they know both the firm and the people in their network. By contrast the "limited choices" hypothesis stresses that finding employment through social networks limits the opportunities and match quality (Loury, 2006). In addition, current employees may have personal interests in referring friends that conflict with the interest of

[12]Research has shown that hiring is a 'stage specific process': factors that affect the initial screening decisions vary from those that influence decisions at the interview stage (Dipboye et al., 1975)

[13]Schoer et al., 2014 report that up to 68% of workers found employment via social networks.

the firm (Fafchamps and Moradi, 2015; Beaman and Magruder, 2012).

A formal referral system with endorsements from former employers may thus be a more effective mechanism to reduce information asymmetries. Interviews with South African firms confirm the benefits of having former employers as references: if available, hiring managers report that they typically call them for the group of shortlisted candidates. However, focus group discussions with job seekers reveal that most do not have contactable references listed on their CV and less than 5% used a reference letter as part of the application process.[14]

1.2.2 Conceptual Framework

Markets differ in the extent to which references can mitigate information asymmetries. In many markets, sellers have no choice over the source of the reference and whether this information is publicly revealed.[15] By contrast, job seekers typically choose referees and often observe their feedback before deciding whether to reveal it to the market.[16] This is an important feature which may limit how effectively referral systems can reduce information asymmetries. This section introduces a static illustrative framework for employer learning in this type of market. It will generate two important sets of results: i) it identifies conditions under which letters have value and ii) it derives predictions for how the letter affects the hiring decision and screening ability of firms.

[14]Furthermore, of those that list references, many include relatives or friends – something the employers in our interviews attached little value to. While reference letters from people in non-professional networks such as former teachers or pastors is sometimes perceived as adding value, employers generally felt that references from former employers provide a much stronger and more credible signal.

[15]The rise of the internet facilitated the exchange of information between former and potential buyers (Avery et al., 1999). Online markets like Amazon and AirBnB provide feedback from former buyers. In online labor markets like odesk, employers are required to publicly evaluate former employers.

[16]Research confirms that reference letters tend to be overly positive (Aamodt et al., 1993). This 'leniency bias' limits the letters' informativeness (Loher et al., 1997)

Setup A job-seeker has (general) ability a which determines her productivity for any firm. At the time of applying for work, the job-seeker is endowed with an application signal $s_1 = a + e_1$. This represents the content of a resume, including school transcripts and other easily observable applicant attributes. With probability π she is also endowed with a reference letter signal $s_2 = a + e_2$ ($c = 1$ if she does, otherwise $c = 0$).[17] Assume that $a \sim nid(0, 1)$, $e_1 \sim nid(0, \sigma_1^2)$ and $e_2 \sim nid(0, \sigma_2^2)$. The job-seeker applies to a vacancy by sending application s_1 to the firm and must choose whether to also attach a reference letter s_2 ($d = 1$ if she does, otherwise $d = 0$).

The firm offers a fixed wage and chooses whether to hire the applicant based on available information Ω. It will do so if the expected productivity exceeds the cost of employment θ, i.e. $E(a|\Omega) > \theta$. We denote this hiring decision as $h = 1$ if a job is offered and $h = 0$ otherwise. Her utility depends only on whether or not she is offered a job, and there is no cost to applying or sending reference letters. The firm's conditional expectation is rational and common knowledge, but the hiring threshold θ is private information.

Solution and Interpretation Applying the perfect Bayesian equilibrium to this dynamic game of incomplete information produces a single stable solution.[18] Although the model outcomes cannot generally be expressed as closed-form solutions of the model parameters, we use linearization techniques to obtain such expressions. (For a formal derivation see Appendix A.)

The firm's equilibrium conditional expectation function depending on whether they receive a reference letter (d) can be expressed as

[17] Building on Gibbons and Katz, 1991, we assume that π is independent of a which limits what firms can infer about workers ability from their access to letters. (Predictions would not qualitatively change as long is there is no perfect correlation.) This assumption is supported by field work we conducted finding that some firms out of principle do not provide reference letters to former employers, citing concerns about legal reasons.

[18] After ruling out the possibility that no-one sends a reference letter, in which case the firm's conditional expectation for those with a reference letter would be undefined.

$$E(a|s_1, s_2, d = 1) = \kappa_2 s_1 + \kappa_1 s_2 \tag{1}$$

$$E(a|s_1, s_2, d = 0) = -\psi\kappa_1\omega + (\frac{1}{1+\sigma_1^2}\kappa_1 + \kappa_2)s_1 \tag{2}$$

where κ_1 and κ_2 capture noise in resumes (s_1) and reference letters (s_2), respectively. ω is the reference variance conditional on the information in the resume and ψ is a monotonic transformation of $P(c = 1|d = 0)$, the number applicants who have a reference letter but choose not to attach it, expressed as a share of all those who do not attach a reference letter.

When applicants include a reference letter ($d = 1$), firms form beliefs about ability using information from both the resume and reference letter, weighted according to the relative reliability of these two signals. If the application does not include a reference ($d = 0$), firms form beliefs using the information in the resume. They further penalize these applicants with a downward adjustment in expected ability, conditional on the quality of the resume. The magnitude of this penalty ($\psi\kappa_1\omega$) increases in the share of applicants who have access to letters, the relative reliability of the letter and the variance of the letter signal.

In equilibrium, applicants with access to letters will choose to send it if it improves the firm's perception of their ability, i.e. $E(a|s_1, s_2, d = 1) > E(a|s_1, s_2, d = 0)$.[19] This requires that the reference is sufficiently positive relative to the information in the resume:

$$d(s_1, s_2,) = c.1\left[s_2 - \frac{1}{1+\sigma_1^2}s_1 > -\frac{0.8\psi}{1-0.64\psi}\omega\right] \tag{3}$$

Predictions Implicit in the model setup are two testable assumptions about the information provided by reference letters: i) letters must be informative about the applicant's

[19]The share of individuals who send reference letters in equilibrium is then: $P(d = 1) = \pi\Phi\frac{\frac{0.8\psi}{1-0.64\psi}\omega}{\tilde{\kappa}_2}$ where $\tilde{\kappa}_2$ is another reflection of the relative reliability of resumes.

ability, i.e. $\frac{\delta E(s_2|a)}{\delta a} > 0$, and ii) letters must contain information that is not already contained in the applicant's resume, i.e. $\frac{\delta E(s_2|a,s_1)}{\delta a} > 0$. Under these assumptions, the model makes the following predictions about how job seekers use letters and how firms respond to receiving letters.

1. *Hiring probability*: Firms will be more likely to hire candidates with stronger letters: $\frac{\partial P(h|s_1,s_2,d=1)}{\partial s_2} > 0$.

2. *Screening on ability*: Letters results in a closer mapping from ability to job offers: $E(a|h=1,d=1) > E(a|h=1,d=0)$.

3. *Credibility*: Since the effect of reference letters depends on their relative reliability (κ_2), any attribute that casts doubt over their reliability (e.g. not providing contact information or being implausibly positive) reduces their effectiveness: $\frac{\partial^2 P(h|s_1,s_2,d=1)}{\partial s_2 \partial \kappa_2} < 0$.

4. *Variance in (prior) beliefs*:

   (a) If employers are more uncertain about ability of job seekers, then the content of reference letters matters more: $\frac{\partial^2 P(h|s_1,s_2,d=1)}{\partial \sigma_1^2 \partial s_2} = \frac{\sigma_2^2}{(\sigma_2^2+\sigma_1^2\sigma_2^2+\sigma_1^2)^2} > 0$.

   (b) Evidence suggests that in our study context, employers are more uncertain about skills of female job seekers. The content of women's letters therefore has a larger effect on the hiring probability: $\frac{\partial P(h|s_1,s_2,d=1,female)}{\partial s_2} > \frac{\partial P(h|s_1,s_2,d=1,male)}{\partial s_2}$.

5. *Usage of letters*:

   (a) As more job seekers gain access to reference letters (π), the usage will increase for two reasons: i) mechanically, more people will have access to positive letters that meet condition 3 and ii) on the margin, people with less positive letters will use it as the penalty of not sending the letter ($\psi\kappa_1\omega$) increases.[20]

[20] A variation of this prediction is known as the Full Disclosure Theorem: if certification of types is costless, then there is full disclosure of information (Grossman and Hart, 1980; Milgrom, 1981).

(b) As access to and usage of letters increase, the ability of firms to identify higher ability candidates improves: $\frac{\partial^2 P(h)}{\partial a \delta \pi} > 0$.

The next section describes the experiments we conduct to test these predictions. Section 1.4 reports results on the value of reference letters (Predictions 1 and 2). Section 1.5 provides evidence on on the role of credibility (Prediction 3) and variance in prior beliefs (Prediction 4), as well as the two testable assumptions regarding the letters' content. Appendix A provides simulation results on how the effects of reference letters change as they become more widely adopted (Prediction 5). This framework presens a rational benchmark model, which assumes that job seekers have correct beliefs about the value of reference letters. We revisit this assumption in Section 1.6.

1.3 Study Design

This section first describes the sampling and the process of eliciting reference letters common to all three experiments. We then describe each of the experimental designs in detail.

1.3.1 Study Sample

Our target population are unemployed youths between the ages of 18 and 34.[21] We limit our study sample to African unemployed job seekers who have some form of previous work experience (as our interventions tests reference letter from previous employers), have not completed university-level tertiary education and live within traveling distance from our four implementing labor centers in the Gauteng and Limpopo province.[22]

[21]Table A.3 provides summary statistics for job seekers in our sample: 50.2% are female and the average age is 27.3 years. The average level of education is 12.1 years and 67% have completed secondary school (matric). 7% of participants are married and they have on average one child. 11.4% receive unemployment insurance and the average participant spends 14 hours per week searching for work.

[22]We worked with the following centers: Krugersdorp, Sandton and Soweto in the Gauteng Province and Polokwane in Limpopo.

Job seekers who meet these criteria were randomly selected from the Employment Services South Africa (ESSA) data base. We further stratify the sample by gender to facilitate subgroup analysis. In the recruitment call, surveyors explain that the job seeker is invited to participate in an employment service study at the local labor center on a specified day. In return, they receive a stipend of 30 Rand (2 USD) to cover travel cost. Across all experiments, 67% of the successfully contacted unemployed individuals agreed to participate.[23]

Obtaining Reference Letters We conducted more than 30 interviews with employers who frequently mentioned the importance of contactable references in the screening process. When asked what information they collect from references, employers listed both non-cognitive skills like motivation, reliability and work ethic as well as cognitive skills like numeracy and literacy. They are also interested in the nature of the relationship between the referee and job seeker and why the employment relationship ended. Based on this feedback, we designed a reference letter template that employers can easily fill out. (Appendix A shows figures of the template and examples of completed reference letters.)

The study employs an encouragement design implemented in cooperation with the Department of Labour (DoL). A baseline survey is administered through an in-person interview at the labor center, followed by an aptitude test that evaluates basic math and literacy skills.[24] Next, participants assigned to the treatment group have a brief individual meeting explaining the benefit of obtaining a reference letter and instructions how to use it in the job search. This is followed by a discussion of the job seekers' work history and identification of potential

[23]Using the limited demographic information provided in ESSA, we find that age and gender are not correlated with the decision to participate. By contrast, every year of additional education increases the probability of participation by 1.6 percentage points (p-value: 0.063). Of those that agree to participate, 63.5% showed up at the labor center on the specified day. None of the socioeconomic variables predict whether participants fail to show up at the agreed time and day.

[24]The test takes about 20 minutes and was designed by the researchers. It closely follows standard entry level tests used in the hiring process by large employers in South Africa. Figure A.4 shows that results are approximately normally distributed with a mean (median) joint numeracy and literacy score of 61% (63%). For sample questions see Figure A.3.

referees. We provide job seekers with several hard copies of the template and instructions on how to return the completed letter to us. After one week, participants receive a text message reminding them to obtain and return the reference letter.

31% of encouraged job seekers returned the completed letter. In surveys after five weeks, 42% of people claim to have obtained a letter. Appendix A investigates which characteristics are correlated with the probability of obtaining a letter. Age is the only statistically significant predictor of receiving a letter; however, there are likely unobservable variables correlated with the propensity to obtain a letter.[25]

Table 1 provides summary statistics of the reference letter content, converting employer ratings into numeric values (0=below average, 1=average, 2=good, 3=very good). Overall, ratings tend to be positive: on a scale from 0 to 6, the average aggregate hard and soft skill rating is 4.9; 11% have a perfect score of 6. We find that hard skills are slightly less positively rated than soft skills (2.3 vs. 2.6 on a 3 point scale). While for most categories women receive slightly more positive ratings, only one gender difference is significant at the 10% level (Team Ability) and one at the 5% level (How highly recommended). As an important caveat, we do not verify the authenticity of the reference letters. In Section 1.5, we will explore whether the letter provides truthful information.

[25]Older job seekers are significantly more likely to have the letter completed, whereas the coefficients of both education and gender are small in magnitude and not statistically significant (Table A.1, Column 1). We find that neither people who are more actively searching nor those who perform better on the aptitude test are more likely to obtain a letter (2,3). The time since last employment is uncorrelated (5) and the employment spell in the last job is negatively correlated (4), although the latter relationship loses significance when we control for all characteristics simultaneously (7). People who were fired in the last job are less likely to obtain a letter than those that left voluntarily, but these differences are not statistically significant (6).

Table 1: Content of Reference Letter by Gender

			Gender		
	N	mean	Female	Male	p-value
Total Score	119	4.933	5.04	4.821	.134
Hard Skill Score	119	2.307	2.362	2.25	.211
Soft Skill	120	2.625	2.677	2.571	.151
All Positive	119	0.109	0.131	0.086	0.434
TeamAbility	117	2.692	2.77	2.607	.058
WorkEthics	120	2.675	2.742	2.603	.162
Reliability	118	2.568	2.597	2.536	.568
Agreeability	118	2.61	2.645	2.571	.448
Interpersonalskills	119	2.597	2.639	2.552	.408
Literacy Ref	117	2.462	2.5	2.421	.487
Numeracy Ref	115	2.174	2.22	2.125	.48
ComputerLiteracy	109	1.917	2.052	1.765	.104
LearningAbility	118	2.576	2.574	2.579	.961
Task1	70	2.5	2.5	2.5	1
Task2	60	2.433	2.452	2.414	.807
Comments (any)	120	.458	.452	.466	.88
Comments (nr)	120	1.842	1.984	1.69	.606
How Recommend (0=reserv.,2=highly)	104	1.558	1.691	1.408	.012
Confidence Assessing (0=low, 2=high)	112	1.67	1.717	1.615	.278
Termination: Voluntary	107	.224	.263	.18	.304
Termination: Contract Ended	107	.645	.632	.66	.762
Termination: Retrenchment	107	.112	.088	.14	.403
Termination: Fired	107	.019	.018	.02	.927
Signed	115	.974	.967	.981	.63
Phone listed	115	.957	.934	.981	.205
Email listed	115	.496	.492	.5	.931

Notes: The table results details from the completed reference letters. Ratings are converted to numeric values (0=below average, 3=very good). Columns on the right provide summary statics separately for women and men and report p-values of a test of equal means.

1.3.2 Experimental Designs

Figure 1 describes how our intervention may affect employment and summarizes our experimental designs. In Experiment 1, we submit applications on behalf of job seekers to vacancies from online job sites and test if employers are more likely to respond if a reference

letter is attached. This provides a "test case" whether reference letters have the potential to be valuable. Experiment 2 explores the effect of reference letters on job search behavior and estimates employment effects *after* people adjusted their search strategy. Experiment 3 tests different forms of encouragement to investigate why only a small share of people obtain reference letters in equilibrium.

Figure 1: Experimental Design Overview

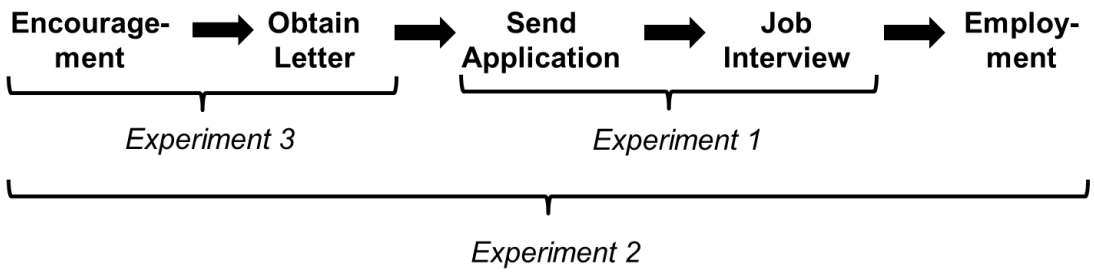

Experiment 1: Employers' Response to Reference Letters To test the effect of the letter on employer demand, we employ a within-subject randomization design: we encourage 441 job seekers across three labor centers (Soweto, Sandton, Krugersdorp) to obtain a reference letter using the protocol described above; for the 31% of participants who return it to us, we send out applications with and without the reference letter.[26] This has the advantage that we can control for individual specific factors that determine employer responses and thus estimate the effect of reference letters more precisely.

Figure 2 summarizes the randomization design. We search the four most popular South African job websites to identify vacancies for entry positions from one of the following sectors: administration, call center, cleaner, driver, retail, security and unskilled. The vacancies are randomly assigned to vacancy slot 1 through 6. Next, we select four of the job seekers who returned the letter and have previous work experience in a related sector. We create email

[26]Selection at the encouragement stage may affect the generalizability of results. However, using within-subject randomization ensures that results are internally valid.

Figure 2: Experiment 1: Randomization Design

Participant \ Vacancy	1	2	3	4	5	6
A	CV + Ref Letter	CV	CV	CV	CV	CV + Ref Letter
B	CV	CV + Ref Letter	CV	CV	CV	CV + Ref Letter
C	CV	CV	CV + Ref Letter	CV	CV	CV + Ref Letter
D	CV	CV	CV	CV + Ref Letter	CV	CV

addresses for each participant and send out six applications following the pattern described in Figure 2. For example, for Participant A we send four applications with the CV (and any additional supporting documents the job seeker provides) and two applications for which we attach the reference letter as an additional document. Importantly, we are invisible to the employer in the entire application process.

Vacancies 1 through 4 offer a straightforward test of the effect of reference letters as we can compare employer responses between applications with and without the attached letter (e.g. compare cell A1 to cell A2, A3 and A4). For vacancy 5 we only send CVs. This provides us with a test for displacement effects at the interview stage, i.e. whether being in an application pool with somebody with a reference letter reduces the chances of getting an employer response. To test for this, we can compare employer responses in cell A5 to A2, A3, and A4. Vacancy 6 receives three applications with reference letters. Comparing application A1 and A6 allows us to test whether employers respond to reference letters differently once they present a higher proportion of the applicant pool.

We submitted a total of 2,050 applications for 102 job seekers between June 2015 and April 2016.[27] We regularly checked for firm responses and forwarded these to the job seekers.[28]

[27]A total of 117 letters were returned to us, of which 15 letters were either illegible or these job seekers did not have work experience in a relevant sector. We included vacancy 6 starting in January 2016.

[28]One possible concern is that employers may contact job seekers directly via phone. From talking to participants this did not happen frequently. While it may lead us to underestimate the overall response rate

Experiment 2: Job Search and Employment Effects While Experiment 1 cleanly identifies the effect of including a reference letter in applications, it does not allow us to test whether people search differently once they obtain a letter. South African job seekers use a mix of search strategies beyond online vacancies (Schoer et al., 2014) and employment effects are more meaningful if they are measured *after* people adjusted both search intensity and search channels. We therefore conduct a second experiment with a separate sample in which half of the job seekers receive the encouragement treatment described in Section 1.3.1.

A total of 1,267 participants are part of this sample and were initially surveyed between September 2015 and February 2016. Participants are invited to come to the labor center on a certain date, randomly assigned to be either control or treatment days. The same calling script is used for the control and treatment group to ensure that there is no differential selection. The share of invited participants who show up are very similar (64.2% reference letter, 63% control group, p-value of test of equal coefficient: 0.55).[29]

To track job search activities and employment outcomes over time, we conduct phone surveys five weeks and three months after the treatment.[30] One potential shortcoming of any survey data is that it is self-reported. We therefore complement the survey data with an observed measure of job search. Specifically, study participants are notified about a vacancy and are asked to submit their full application via email in case they are interested.[31]

---

there is little reason to believe that the choice of how employer communicate with job seekers is correlated with the treatment assignment.

[29]Table A.3 suggests that the randomization was successful. Of the 19 characteristics reported, only two differences are significant at the 10 percent level. When we control for all characteristics in a regression, none of the variables is significant and we can reject that variables are jointly significant (p-value: 0.72, results not reported).

[30]Table A.4 shows that attrition rate increases from about 6% in wave 1 to 17% in wave 2, likely due to survey fatigue and participants switching phone numbers. Attrition is clearly not random: younger and less educated participants are more likely to attrite, but importantly rates do not differ between treatment and control group.

[31]Participants were informed about a vacancy in a specific sector. Among those with work experience in multiple sectors, we randomly chose for which sector we notify them. For job seekers for who we do not have information on previous sectors, we send a general notification about a vacancy. Sectoral shares were

Experiment 3: Barriers to Obtaining Letters Results discussed in more detail below suggest that reference letters substantially increase the probability of receiving an employer response. This raises the question of why only about 2% of job seekers in the control group use reference letters in their job search. Experiment 3 tests different barriers to obtaining reference letters.

During follow up surveys, a significant share of participants could not provide us with a reason why they have not tried to obtain the letter or cited reasons like "No Time" or that they do not need it. This may be a sign of procrastination or that job seekers do not believe they would benefit from a letter. We design two interventions to test these hypotheses: i) provide job seekers with information on the effectiveness of letters and ii) compensate participants with 100 Rand (about half a daily wage) in cell phone airtime if they obtain a letter.[32]

A group of 498 job seekers, previously encouraged to obtain a letter, receives a follow-up text message to their cell phone and (if provided) email address reminding them of how to return the completed letter to us. Participants were randomized into four groups.[33] The control group received only this reminder. The other three groups received one of the following additional messages:

- "*Research suggests reference letters almost double chances of getting a job interview.*" (Information)
- "*To compensate your costs, you get 100 Rand airtime after sending us the completed letter.*" (Compensation)

balanced by treatment status. Applications were submitted to actual vacancies after the end of the last survey wave so that it would not confound employment estimates.

[32]Job seekers were compensated with airtime rather than money for logistical reasons. Participants could provide any time for the airtime top up and could choose among all major cellphone carriers. As participants frequently top up cell phone airtime, we expect that they value it similar to receiving 100 Rand in cash.

[33]Comparing observable characteristics between the treatment and control group suggests that randomization was successful (Table A.5).

- "*Research suggests reference letters almost double chances of getting a job interview. To compensate your costs, you get 100 Rand airtime after sending us the completed letter.*" (Information + Compensation)

1.4 Do Reference Letters Have Value?

1.4.1 Empirical Strategy

This section tests the effect of the reference letter on firm demand using data from Experiment 1. We use two measures of employer response: i) a narrow measure of interest that captures interview requests and ii) a broader measure of interest that captures either an interview request or a different employer response (most commonly, firms asked questions, requested specific documents, or provided more information about the job and asked if job seekers were still interested). Throughout the analysis we will report results for both outcomes.

To estimate the effect of the reference letter, we estimate the following model:

$$y_{is} = \beta Ref_i + \lambda_s + \mu_k + e_s \tag{4}$$

Outcome y_{is} is a binary variable measuring whether employers respond to application i of person s. Ref_i is an indicator variable for whether a reference letter was included with application i. λ_s and μ_k capture individual and sector fixed effects, respectively. The error term e_s is clustered at the individual level. The coefficient of interest β captures the causal effect of the reference letters.

1.4.2 Employer Responses

Table 2 reports results from Specification 4. Column 1 to 4 report effects using the broad measure of interest as an outcome and Column 5 to 8 report effects on interview requests. The preferred specification controls for both sector and individual fixed effects (Column 3, 7). On a control mean of 4.15 percent, the reference letter increases the chance of getting any employer response by 2.44 percentage points (59%) (3) and on getting an interview request by 1.44 percentage points (60%) on a control mean of 2.4 percent (7). The former effect is significant at the 5% level, the latter is not significant at conventional levels (p-value: 0.11).[34] Treatment coefficients for both outcomes are substantially larger for women, although the difference across gender is not statistically significant (4, 8).[35]

Table 2: Effect of Reference Letter on Call Back

	y=Employer Response: **Interest**				y=Employer Response: **Interview**			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Reference Letter	0.0254**	0.0251**	0.0244**	0.0105	0.0154*	0.0147*	0.0144	0.0037
	(0.0102)	(0.0102)	(0.0107)	(0.0162)	(0.0087)	(0.0087)	(0.0091)	(0.0115)
Female				-0.0107				-0.0037
				(0.0141)				(0.0085)
Female x Letter				0.0268				0.0203
				(0.0211)				(0.0163)
Sector F.E.	N	Y	Y	Y	N	Y	Y	Y
Individual F.E.	N	N	Y	N	N	N	Y	N
R^2	0.003	0.01	0.078	0.009	0.002	0.016	0.057	0.010
N	2050	2050	2050	2050	2050	2050	2050	2050
Control mean	0.0415	0.0415	0.0415	0.0415	0.0240	0.0240	0.0240	0.0240

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (reported in parentheses) are clustered at the applicant level. Results report OLS estimates. Dependent variables are binary measures of employer response: interview requests (Col. 4-6) and either interview request or a different employer response expressing interest in the job applicant (Col 1-3). Sector fixed effects are included for the six sectors for which we send applications.

Do employers respond differently if they receive multiple applications with reference letters? We estimate Specification 4 including an interaction term between the reference letter vari-

[34]Table A.6 reports results of Specification 4 estimated separately for each of the seven sectors. There is some evidence that effects are larger in sectors that require fewer cognitive skills (e.g. cleaner, retail, unskilled). However, the evidence is inconclusive given the relatively small number of applications sent to each sector.

[35]Within the sample of women, we find significant effects for the interest outcome (at the 1% level) and interview outcome (at the 5% level). Treatment coefficients in the sample of men are not significant.

Table 3: Multiple Reference Letter and Displacement

	Y=Interest			Y=Interview		
	(1)	(2)	(3)	(4)	(5)	(6)
Reference Letter	0.02429**	0.02364**	0.02305**	0.01437	0.01404	0.01334
	(0.0110)	(0.0117)	(0.0114)	(0.0090)	(0.0099)	(0.0094)
Reference Letter x Multiple		0.0044			0.0023	
		(0.0305)			(0.0254)	
Control Group - Pure			-0.00827			-0.00689
			(0.0127)			(0.0103)
R^2	0.080	0.080	0.080	0.058	0.058	0.058
N	2050	2050	2050	2050	2050	2050
Control mean	0.0415	0.0415	0.0415	0.024	0.024	0.024

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at applicant level.
Coefficients report results of Specification 4. Column 2 and 4 include an interaction term between the reference letter indicator and an indicator of the vacancy that receive three reference letters. Column 3 and 5 includes a dummy for applications sent to a vacancy that does not receive any reference letters.

able and an indicator variable for vacancy 6, which received three applications with letters. The coefficient on the interaction term is very close to zero indicating that the effect does not differ if the employer receives more than one letter (Table 3, Column 2, 5). These results suggest that it is not the novelty of seeing a reference letter that is driving the positive employer response.[36]

Next, we test if there is a negative effect from being in the applicant pool with a job seeker who submits a reference letter. We include a dummy for pure control applications (sent to vacancy 5) in Specification 4. Coefficients in Table 3 are small in magnitude and not statistically significant suggesting that there is no displacement (3, 6). However, these coefficients are estimated relatively imprecisely and we can only rule out with 95% confidence that the displacement effect is larger than 1.2 percentage points.

[36]However, the exact interpretation remains inconclusive as seeing multiple letters that use the same template may increase its credibility. It further allows employers to compare letters and calibrate the ratings.

1.4.3 Screening Ability

The starting premise of the paper is that information asymmetries inhibit firms to identify the most suitable candidates. Following the model in Section 1.2.2, we assume that there is an ability parameter a, imperfectly observed by the firm at the time of the application.[37] As a proxy for productivity, we employ standardized results of the aptitude test administered as part of the baseline survey.

To test whether the letters enable firms to identify applicants of high ability (Prediction 2) we estimate the following model:

$$y_{is} = \beta Ref_i + \gamma a_s + \delta Ref_i * a_s + \mu_k + e_s \tag{5}$$

Table 4: Effect of Reference Letter on Screening Productive Applicants

	Y=Interest		Y=Interview	
	(1)	(2)	(3)	(4)
Reference Letter	0.02575**	0.00838	0.01555*	0.00522
	(0.010)	(0.016)	(0.009)	(0.012)
Aptitude (z-score)	0.00618	0.00801	0.00062	0.00047
	(0.005)	(0.009)	(0.003)	(0.004)
Ref Let x Aptitude (z-score)	0.01999**	0.01574	0.01305**	0.01230
	(0.008)	(0.011)	(0.006)	(0.008)
Ref Let x Female		0.03166		0.02078
		(0.022)		(0.018)
Ref Let x Female x Aptitude (z-score)		-0.00271		-0.00807
		(0.011)		(0.015)
R^2	0.003	0.008	0.002	0.004
N	2050	2050	2050	2050
Control mean	0.0415	0.0415	0.0240	0.0240

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (reported in parentheses) are clustered at the applicant level. Results report OLS estimates controlling for sector fixed effects. *Aptitude* is measuring the standardized English and Math score. For readability reasons, we suppressed coefficients for *Female* and *Female x Aptitude*. These coefficients are small in magnitude and insignificant.

[37] This parameter captures the part of productivity that is transferable across firms rather than firm specific match quality.

Coefficient γ captures whether employers select higher ability applicants when only the CV is attached and $\gamma + \delta$ is the effect when the letter is attached. Results are presented in Table 4. It is noteworthy that coefficients γ are small in magnitude and not significant suggesting that without the reference letter, firms are ineffective in selecting the more productive applicants. δ is positive indicating that reference letters enable firms to identify applicants of higher ability (despite not seeing the aptitude score). The coefficients are significant at the 5% level and large in magnitude. A one standard deviation higher performance in the aptitude test increases the probability of receiving an employer response and interview request by 2 percentage points (47%) and 1.3 percentage points (54%), respectively (Column 1, 3). Put differently, in control applications the chance of receiving an employer response for job seekers at the 90th ability percentile is 1.8 percentage points (35%) higher compared to those at the 10th percentile. Once the reference letter is included this figure increases to 6.3 percentage points (123%). These improvements in firms' screening ability does not differ by the gender of the job seeker (2, 4).

1.4.4 Employment Effects

To test whether reference letters are increasing firm responses and employment when used by job seekers, we use data from Experiment 2 and estimate the following model:

$$y_{ij} = \beta T_i + \gamma X_i + \delta y_{ij}^{bs} + \lambda_j + e_i \tag{6}$$

The dependent variable y_{ij} is measured for individual i residing in location j. We focus on three key outcomes: number of applications submitted and job interviews in the last four weeks and employment status. In order to increase precision we control for the baseline value y_{ij}^{bs} of outcomes. To account for differences in firm demand across space, we control for location fixed effects λ_j. Robust standard errors are computed at the individual level.

Results from the audit study (Table 2) suggest that the effect of reference letters may differ by gender. We therefore also estimate specification 6 separately for women and men.

Table 5: Effect of Reference Letter on Employment (3 months)

	Intent to Treat Effects			**Local Average Treatment Effects**		
	(1)	(2)	(3)	(4)	(5)	(6)
	Application	Interview	Employment	Application	Interview	Employment
			Panel A: POOLED			
Reference Letter	0.660	0.072	0.020	1.336	0.147	0.037
	(0.426)	(0.046)	(0.022)	(0.857)	(0.092)	(0.046)
R^2	0.222	0.051	0.015	0.222	0.046	0.008
N	997	996	1033	997	996	1033
Control Mean	3.975	0.675	0.130	3.975	0.675	0.130
			Panel B: FEMALE			
Reference Letter	1.051	0.130**	0.057*	2.249	0.280**	0.117*
	(0.702)	(0.059)	(0.032)	(1.522)	(0.125)	(0.068)
R^2	0.267	0.063	0.029	0.242	0.050	0.001
N	501	506	528	501	506	528
Control Mean	3.842	0.534	0.117	3.842	0.534	0.117
			Panel C: MALE			
Reference Letter	0.118	0.014	-0.015	0.553	0.027	-0.032
	(0.431)	(0.071)	(0.032)	(0.868)	(0.135)	(0.062)
R^2	0.282	0.042	0.021	0.232	0.041	0.020
N	491	492	510	491	492	510
Control Mean	4.130	0.862	0.157	4.130	0.862	0.157
p-value: $\beta_{fem} = \beta_{male}$	0.368	0.241	0.090			

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Results presented in Column 1-3 are intent to treat estimates. Results in Column 4-6 are treatment on the treated estimates, using the encouragement assignment as an instrument for take-up. All regressions control for covariates. Panel A reports estimates from Specification 6 for the full sample. *Application* and *Interviews* measures the number of applications submitted and job interviews in the last four weeks, respectively. The number of applications and interviews are winsorized at the 1% level to account for outliers. Employment is an indicator variable denoting if people are in paid employment or self-employed. Panel B and C estimate results separately for women and men.

Columns 1 to 3 in Table 5 report intent to treat (ITT) effects after three months. Columns 4 to 6 reports local average treatment effects (LATE) estimated with 2SLS, using the random encouragement assignment to instrument for the take-up of reference letters.[38] Results in the

[38]The first stage equation is $Ref_i = \gamma T + \beta X + e_i$ with T_i as the randomly assigned encouragement treatment. Local average treatment effects for the group of compliers. ITT and LATE estimates are both interesting from a policy perspective. ITT estimates show what effect we can expect from scaling-up programs

pooled sample are inconclusive (Panel A): coefficients on both the number of applications submitted and on employment outcomes are sizable - LATE estimates range between 20% and 30% of the control mean - but not statistically significant.

Panel B and C show that there is important treatment effect heterogeneity by gender: after three months, women in the treatment group submit more applications and are significantly more likely to receive interviews and find employment. Employment effects are large in magnitude: 5.7 percentage points for ITT estimates (3) and 11.7 p.p. for LATE estimates (6), effectively doubling employment rates for the group of compliers. Coefficients for men are close to zero and insignificant. We can reject that employment coefficients for women and men are equal at the 10 percent level. While estimates are relatively imprecise, these results suggest that reference letters have the potential to improve employment outcomes.[39]

Results in this section provide support for model Prediction 1 and 2. Next, we explore the mechanism underlying these large and significant effects of reference letters.

1.5 How Do Firms Use Reference Letters?

1.5.1 Are letters informative?

This section tests the two key assumptions necessary for reference letters to be effective: they must be informative of applicants' skills and provide information that cannot easily be inferred from other application documents. We test these assumptions by comparing subjective employer ratings to an objective assessment. Specifically, we regress results from

---

as this measure accounts for the fact that some job-seekers will fail to take advantage of the service. LATE estimates, by contrast, provide information about the actual effect of a program for the compliers (Angrist and Imbens 1994).

[39]Table A.8 reports results after five weeks. Results are smaller and insignificant, possibly because the follow up period is too short as many participants report that it takes them longer to obtain a reference letter.

the aptitude test we administer on employer ratings in numeracy and literacy. Table 6 shows that employer ratings and test results are highly correlated for both literacy (1) and numeracy (4). This implies that the average letter contains information about the applicant's skills. Next, we explore how the correlation changes when we control for additional covariates (age, education, gender) and school grades in English and math, respectively. While the magnitude of the coefficients decreases they stay highly significant suggesting that the letter contains information that employers cannot easily infer from the resume (2, 5).[40] Results do not differ by gender, ruling out that treatment effects are larger for women because their reference letters are more informative (3, 6).

Table 6: Are Numeracy and Literacy Employer Ratings Correlated with Aptitude?

	Literacy: Reference Letter (z-score)			Numeracy: Reference Letter (z-score)		
	(1)	(2)	(3)	(4)	(5)	(6)
Literacy: Aptitude	0.3645***	0.2274**	0.2458**			
(z-score)	(0.0935)	(0.1026)	(0.1185)			
Female x Literacy Apt			-0.04907			
(z-score)			(0.2066)			
Numeracy: Aptitude				0.3001***	0.2627***	0.25585*
(z-score)				(0.0885)	(0.0966)	(0.1381)
Female x Numeracy Apt						0.01548
(z-score)						(0.1788)
Covariate	N	Y	Y	N	Y	Y
School Grade	N	Y	Y	N	Y	Y
R^2	0.136	0.232	0.232	0.093	0.116	0.116
N	116	116	116	114	114	114

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is the standardized value of the numeric employer rating (0=below average, 3=very good). *Literacy* and *Numeracy* measure the standardized performance in the aptitude test. Control variables include age, gender and education. School grade is measuring the grade (in %) participants achieved in the last math and English class, respectively.

Results confirm that referee ratings contain additional information, at least for skills captured in the aptitude test. Arguably, it would even be more difficult for firms to learn about other skills from the CV, especially non-cognitive skills like reliability or work ethics (Aamodt,

[40] Groh et al., 2014 employ a similar test on a sample of unemployed youths in Jordan and find that results from psychometric and skill tests have predictive power for subsequent employment, even after controlling for easily observable worker characteristics. Abebe et al., 2016 find that the job search workshop improves firms' ability to identify applicants' whose observable characteristics predict higher performance in aptitude tests.

2015).

1.5.2 Does the Reference Letter Content Matter?

The model predicts that those with *better* reference letters are more likely to receive job offers (Prediction 1). An alternative channel is that the mere ability to obtain a reference letter and use it correctly in the job search is the relevant signal for firms. We can test for whether employer responses depend on the content of the letter, by estimating:

$$y_{is} = \mu_k + \lambda_s + \beta Ref_i + \gamma score_s + \delta Ref_i * score_s + e_s \tag{7}$$

Coefficient γ captures the counterfactual, i.e. the effect of the referee rating (*score*) when it is not revealed to employers, indicating whether job seeker that are in higher demand receive more positive reference letters. We find that the coefficient is close to 0 (Table 7). Coefficient δ measures the (additional) effect of the referee rating once the letter is revealed to the firm. The score is positive but not significant for both outcomes (1, 5). Looking at the relationship between referee ratings and employer responses graphically shows a non-linear relationship: ratings and employer responses are positively correlated, but we observe a sharp discontinuity for letters with perfect scores. We therefore estimate specification 7 and control for applications with perfect scores (3, 7). The coefficient on the rating increases and turns significant: a one standard deviation higher rating increases employer responses (interviews) by 41% (71%). The coefficient on the all positive dummy interacted with the reference letter is negative and large in magnitude, but only significant (at the 1% level) for the interview outcome (3, 7). A letter with a perfect score has a 7 percentage points lower chance of receiving an interview compared to what is predicted by the rating (7).

The content of the letter matters much more for female applicants: positive ratings have a larger positive impact and letters with perfect ratings have a more negative effect across

Table 7: Effect of Referee Rating on Call Back

	Y=Interest				Y=Interview			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Reference Letter	0.0371	-0.0014	0.0441	-0.0033	0.0447	0.0268	0.0531	0.0156
	(0.0561)	(0.0605)	(0.0544)	(0.0584)	(0.0469)	(0.0519)	(0.0454)	(0.0502)
Referee Rating (z-score)	-0.0006	0.0101	-0.0030	0.0083	0.0014	0.0044	0.0007	0.0060
	(0.0057)	(0.0084)	(0.0056)	(0.0065)	(0.0040)	(0.0057)	(0.0045)	(0.0060)
Letter x Rating (z-score)	0.0077	-0.0120	0.0167*	-0.0065	0.0057	0.0002	0.0169*	0.0009
	(0.0086)	(0.0127)	(0.0089)	(0.0117)	(0.0080)	(0.0084)	(0.0092)	(0.0091)
Letter x Rating x Female		0.0396**		0.0558***		0.0089		0.0368**
		(0.0184)		(0.0188)		(0.0140)		(0.0162)
All positive			0.0164	0.0211			0.0048	-0.0288**
			(0.0266)	(0.0766)			(0.0140)	(0.0131)
Letter x All positive			-0.0584	-0.0708			-0.0731***	0.0016
			(0.0353)	(0.0809)			(0.0248)	(0.0183)
Letter x All pos.x Female				-0.0213				-0.1265***
				(0.0897)				(0.0358)
R^2	0.014	0.016	0.015	0.019	0.014	0.015	0.017	0.021
N	2050	2050	2050	2050	2050	2050	2050	2050
Control content	Y	Y	Y	Y	Y	Y	Y	Y
Control mean	0.0415	0.0415	0.0415	0.0415	0.0240	0.0240	0.0240	0.0240

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (reported in parentheses) are clustered at the applicant level. *Total Score* measures the average employer rating converted to numeric values (out of 6). *All positive* is a indicator variable for whether employers give a perfect rating. We estimate with model with all interaction terms but suppress coefficients for readability reasons. All columns control for other content revealed in the reference letter - for detailed results see Table A.7. We include dummy variables for five reference letters that did not include a rating.

both outcomes (2, 4, 6, 8). This finding is consistent with the starting premise that firms are more uncertain about applications from women and thus use letters more for updating beliefs about female job seekers (Prediction 4b).[41]

The discontinuity in employer responses at high scores raises the question of whether firms are correct in inferring that these applicants are of lower ability. These job seekers are in fact the group that performs *best* in the aptitude test.[42] In line with Prediction 3 of the theoretical model, this suggests that employers ignore the reference letter signal if it is perceived to be implausibly positive and thus deemed non-credible. After all, it is unclear why the firm did

[41]These results are consistent with earlier work showing that reducing information asymmetries leads to a larger belief updating among employers for members of disadvantaged groups (Agrawal et al., 2013; Lang and Manove, 2011).

[42]Writing implausibly good reference letters presents a form of inadvertent signal jamming. Results (not reported) confirm that the effect of reference letters on firms' ability to pick higher ability applicants is *increased* when we estimate Specification 5 without these positive letters.

not continue to hire these job seekers if they are "very good" at every skill.[43]

These findings provide empirical support for studies that explore the role of credibility of signals. Clark and Martorell, 2014 conclude that in addition to providing relevant information, signals must be verifiable in order to be of value.[44] Avery and Meyer, 2011 echo this argument and observe that there is no universal standard for the assessment procedure, nor databases on the history of past recommendations. This induces evaluators to be biased thus reducing their usefulness in the hiring process (Avery and Meyer, 2011).

1.6 Why are Reference Letters Not More Widely Used?

The previous section showed that both job seekers and firms benefit from reference letters: they increase workers' chances of receiving callbacks and help employers to pick job seekers of higher ability. This raises the question of why the market is in a near pooling equilibrium in which reference letters are almost completely absent. The previous analysis rules out two of the most obvious explanations, confirming that reference letters contain additional information and, despite being sent by job seekers, employers use them to update beliefs. This section explores additional explanations on the part of previous employer, hiring firms and job seekers.

1.6.1 The Role of Previous Employers, Hiring Firms and Job Seekers

Previous Employers We ask job seekers to bring all their application documents to the initial meeting at the labor center. We find that among job seekers with previous work

[43] A uniform rating may also indicate that the referee did not take the time to carefully consider each category. However, we do not find that the effect of these uniform assessments differs for letters that include more detailed comments on skill categories, suggesting that the negative effect is not due to a perceived lack of effort of the referee.

[44] They speculate that the low return to a high school degree they find may be because diploma information is difficult to verify, partly because schools are under no legal obligation to cooperate with employer requests Clark and Martorell, 2014.

experience, only about 4% have a reference letter from a former employer. When probed, 86.4% of job seekers report that they "Did not ask", while only 3.1% report that they asked but the employer refused (Table 8).

Table 8: Reasons for Low Prevalence of Reference Letters

	N	Mean
Why do you not have a letter? (Baseline)		
I did not ask	936	0.864
Employer refused	936	0.031
It was not requested	936	0.016
Other	936	0.089
Did you try to obtain a letter? (After encouragement)		
Yes	618	0.56
If No, *Why did you not try?*		
Travel Cost / Distance	618	0.052
Firm Unavailable / Relocated	618	0.038
No Time	618	0.037
Bad Terms wit Employer	618	0.019
No Need for it	618	0.013
Other	618	0.281
Did you Succeed? (If participant tried)		
Yes	360	0.736
If No, *Why Not?*		
Firm relocated / unavailable	360	0.078
Waiting to hear back	360	0.053
Firm Refused	360	0.041
Other	360	0.087

Note: Results report responses at different points in time. The first panel asks why participants do not have letters at the time of the baseline. The second panel reports follow up survey responses in the treatment group that was encouraged to obtain a letter. The third panel limits responses to participants that tried to obtain a letter.

It is however possible that many job seekers did not ask because they correctly predict that employers would not be willing to write a letter. We can exploit results from our encouragement design to test this hypothesis. Five weeks after the treatment, 56% of job seekers report that they have tried to obtain a letter. Of this group, 73.6% succeeded. Among those that tried, only 4.1% report that they failed to obtain a letter because the employer refused.

Hiring Firm Results in Section 1.4 indicate that firms believe letters (unless they are implausibly positive), use them to update beliefs of job seekers, and benefit as it enables them to select people of higher ability. Interviews with hiring managers further shows that they recognize that job seekers do not have any bargaining power to request letters. Firms therefore do not require applicants to submit letters.

Job Seeker Why do job seekers not request reference letters from employers? This section reports results from Experiment 3 in which we test the relative importance of the cost and perceived benefits of obtaining letters. We estimate the following specification:

$$y_{ij} = \beta T_i + \gamma X_i + \lambda_j + e_i \tag{8}$$

The outcome y_{ij} is a binary measure of whether individual i residing in location j returned the reference letter. We report estimates with and without controlling for covariate vector X_i. To account for differences across space, we control for location fixed effects λ_j. Robust standard errors are computed at the individual level.

Pooling the information and compensation treatment groups, we find a statistically significant increase in the share of people who obtain a letter of 7.6 percentage points (p.p.) (Table 9, Column 2). This is a sizable effect given the control mean of 18.9 percent. Next, we estimate the effect of each treatment arm separately. The effect of the information treatment is 12.6 p.p. and statistically significant (column 4). By contrast, the effect of monetary incentive is much smaller (2 p.p.) and statistically indistinguishable from 0.[45] We can reject that

[45]Interpreting the small and non-significant effect of the monetary incentive on obtaining letter is not straightforward. Given that participants were compensated in airtime and only after returning the letter, this treatment does not provide a test of liquidity constraints. Instead, the treatment provides a short-term benefit of exerting effort and may thus address procrastinating behavior. There are various potential reasons for why this intervention did not have an effect: the amount may have been too small, the monetary incentive may have crowded out job seekers' intrinsic motivation (Frey and Jegen, 2001) or participants may have regarded the payout to be too far in the future.

Table 9: Take up Experiment

	(1)	(2)	(3)	(4)
Pooled Treatment	0.070*	0.076**		
	(0.0382)	(0.0382)		
Information			0.115**	0.126**
			(0.0513)	(0.0511)
Money			0.008	0.020
			(0.0418)	(0.0422)
Information + Money			0.067	0.066
			(0.0492)	(0.0500)
R^2	0.135	0.152	0.141	0.159
N	499	496	499	496
Control Mean	0.189	0.189	0.189	0.189
Control Variables	N	Y	N	Y
p-value: $\beta_{Inf} = \beta_{Mon}$			0.033	0.038
p-value: $\beta_{Inf} = \beta_{Mon+Inf}$			0.378	0.278
p-value: β_{Mon}=$\beta_{Mon+Inf}$			0.211	0.354

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Column 1 and 2 pool Information, Money and Information+Money groups. The control group received message reminding them of how to return the letter.

treatment effects are identical at the five percent significance level. Interestingly, the effect of providing both information and a monetary incentive (6.6 p.p.) is smaller in magnitude than providing information alone, although we cannot reject that coefficients are identical (p-value: 0.278). The effect of providing information does not differ by gender.[46]

1.6.2 Why do Job Seekers Underestimate Benefits of Letters?

The previous section finds that providing information on the benefits of letters has a large effect on the behavior of job seekers suggesting that erroneous beliefs is one of the reasons why the market is in a pooling equilibrium. This is puzzling as standard learning models predict that job seekers should learn about the returns to having a letter. This section explores why these beliefs can be sustained by investigating how job seekers use letters and

[46]Results (not reported) show that the effects are larger for men (insignificant), which is unsurprising given that we provided information on the effectiveness of letters across gender. Men may thus have received an overly optimistic view of reference letters, whereas women may still underestimate the letters' value.

by comparing how our reference template differs from existing letters.

Low Usage of Letters We first explore results on the usage of the letter from Experiment 2. Participants were informed about an open vacancy and asked to submit their application material if interested. We estimate the following specification:

$$y_{ij} = \beta T_i + \gamma X_i + \lambda_j + e_i \tag{9}$$

using two outcome measures: i) a dummy capturing whether a job seeker i residing in location j submits an application and ii) a dummy measuring whether they submit a reference letter as part of the application. T_i captures whether participants were assigned to the treatment group that received the encouragement to obtain a letter.

Table 10 shows that participants in the treatment group are not more likely to submit applications (1). Next, we investigate the application documents of those that send an application. Unsurprisingly, the share who submits a reference letter is significantly larger in the reference letter group (3). In the control group only 1.1% submit a letter confirming that reference letters are almost completely absent in the labor market we investigate. This share increases in the treatment group: 8% of all participants (or 18.2% of those who obtained a letter) submit it as part of the application (3). While this difference is statistically significant, this figure is far below the share of job seekers who report in the survey to have successfully obtained a letter (44%) and use it in the job search (37%).[47]

We observe a large differences in the usage of reference letter across gender: women are much more likely than men to attach it as part of the application (4, 5). This can in part explain the large difference in employment effects across gender we find in Experiment 2.

[47] One reason is that we asked job seekers to submit material via email and some participants may not have had access to scanners. A larger share of job seekers may indeed use the letter in conventional job search channels.

Table 10: Application Material Submitted

	Y=Submit Application		Y=Attach Reference Letter		
	(1)	(2)	(3)	(4)	(5)
Reference Letter	-0.001	-0.023	0.069**	0.007	-0.000
	(0.022)	(0.033)	(0.029)	(0.030)	(0.006)
Female	0.011	-0.017	0.038	-0.018	-0.003
	(0.023)	(0.032)	(0.029)	(0.021)	(0.004)
Ref Let x Female		0.047		0.113*	0.021**
		(0.045)		(0.058)	(0.010)
R^2	0.006	0.017	0.072	0.091	0.014
N	1141	1141	184	184	1141
Control Mean	0.163	0.163	0.011	0.011	0.002
Sample	full	full	application	application	full

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (reported in parentheses) are clustered at the applicant level. Outcomes are binary measures of whether job seekers submit an application (1-2) and whether they attach a reference letter (3-5). Column 3 and 4 restrict the sample to job seekers who submit an application.

Results suggest that low usage of the reference letters stifles the feedback that job seekers receive about reference letters from the market. The resulting failure to learn is further compounded by the overall low level of search activity among job seekers. Results in Section 1.4 show that the letter reduces the number of applications needed to obtain an employer response from 25 to 15 and an interview request from 40 to 28. However, at baseline the average job seeker only submits about 4 applications per month.

Existing Reference Letters Are Less Informative A second reason why job seekers may underestimate the potential benefit of reference letters is that the type of letter in circulation at the time of the baseline is in fact of lower value. Our model predicts that the effectiveness depends on the noisiness of the reference letter relative to the resume. Reviewing a total of 30 reference letters collected from job seekers in our sample at the time of baseline provides strong support for this hypothesis: the majority of letters lack information on the workers' position (48% include this information), responsibilities (38%), skills (28%) duration of employment (48%), and reason for termination of employment (18%). In addition, only

48% of letters are signed and 56% provide contact information. If job seekers experiment with reference letters that are both less informative and credible, they may incorrectly infer that all letters are ineffective.

In-depth interviews with a sample of 28 hiring firms provide further support for this explanation. 73% of hiring managers report that our reference letter template is more effective than other reference letters they receive. The most frequently cited reasons are that the template provides information on specific skills (55%) and that it is more clearly structured (32%).[48] Asked for reasons that make the template *less* effective than other letters, managers point to the lack of a firm letter head or stamp (45%) and that letters are too positive (14%). This corroborates our experimental findings documenting the importance of credibility and suggests that modifications in the design of the letter may further increase its effectiveness.

1.7 Conclusion

The internet has drastically reduced information asymmetries across many markets: online labor market require firms to provide public evaluations of employees' performance and offer workers to take tests to certify their skills. Services like LinkedIn offer an easy way to communicate credentials, work experience and even endorsements from former coworkers and employers. These professional network sites also identify common connections than can serve as informal references. Yet, large parts of the global labor force is working in markets that have not been affected by these changes.

Our study investigates the role of information asymmetries in one such market: the low-skill sector in South Africa. We document that information asymmetries are prevalent in this

[48] In addition, the rubric form offers *less ambiguous* presentation of the assessment than a reference letter in paragraph form. This may particularly benefit women as previous research documents that candidates who are perceived to be similar by the predominantly male hiring managers receive more favorable evaluations (Cardy and Dobbins, 1986).

market and employers struggle to identify high ability job seekers. We find that a simple intervention - encouraging job seekers to obtain a standardized reference letter from a former employer - can lead to substantial improvements in firms' ability to select job seekers of higher ability from the large pool of applicants. Especially women, who are excluded from many informal referral networks in South Africa, benefit from reference letters. This demonstrates that reducing information asymmetries can improve equity on labor markets.

While our study looks at the effects of reference letter in a static framework, reducing information asymmetries may also have dynamic effects. Similar to many other low-income countries, South Africa suffers from low quality of education, which limits the use of educational credentials to screen job seekers. This has adverse dynamic effects: if a high school degree loses its signaling value, youths may be less motivated to study or graduate. Likewise, if workers are employed on temporary contracts and their job performance is not revealed to the market, returns to exerting effort are lower. Our results suggest that reference letters have the potential to provide a powerful incentive to workers. Reducing information asymmetries may therefore have positive effects on productivity beyond the diminishment of frictions in the matching process explored in this study. Yet, not having a reference letter may also pose a barrier for new labor market entrants as letters enhance firms' ability to screen applicants with job experience relative to entrants.[49] This could lead to inefficiently low hiring of people without work experience (Pallais, 2014). Quantifying these dynamic effects remains the work of future research.

[49]The severity of this barrier depends partly on whether firms are able to distinguish people with and without work experience. If firms do not observe the employment history, they will apply a uniform penalty to applicants without reference letters. By contrast, the penalty is smaller for new labor market entrants if firms can identify applicant types, e.g. from information on the resume.

2 Bridging the Intention-Behavior Gap? The Effect of Plan-Making Prompts on Job Search and Employment

2.1 Introduction

Job search is a complex undertaking.[50] Determining which search channels are likely to be most effective and choosing how much time to invest in alternative search activities can be overwhelming. Often these decisions must be made with little or no feedback on how to increase the probability of receiving a job offer, which can result in sub-optimal search behavior.

As job search is largely self-regulated, there is a variety of psychological and behavioral challenges seekers face. Spinnewijn, 2015, for instance, shows that biases in beliefs about returns to search effort can lead to sub-optimal search intensity. Search efforts have also been shown to depend on job seekers' present-bias (DellaVigna and Paserman, 2005), their locus of control (Caliendo et al., 2015; McGee and McGee, 2015), as well as on self-confidence and willpower (Falk et al., 2006).

The focus of the present paper is on the so-called "intention-behavior gap", defined as the imperfect relationship between the intention to perform a particular behavior and the actual enactment. While the exact mechanisms by which planning prompts promote follow-through are difficult to pin down, there is convincing evidence that planning and scheduling tasks have the potential to help people follow through on behaviors ranging from voting (Nickerson and Rogers, 2010), to getting vaccinations (Milkman et al., 2011), medical screening (Milkman

[50] This paper is co-authored with Rulof Burger (Stellenbosch University), Eliana Carranza (World Bank) and Patrizio Piraino (University of Cape Town).

et al., 2013) and exercising (Prestwich et al., 2003) (see Rogers et al., 2015 and Hagger and Luszczynska, 2014 for recent reviews).

Building on this research, we use a field experiment with a sample of 1100 unemployed South African youth to test the effect of plan-making prompts on search behavior and employment. As part of a job counseling workshop conducted by the South African Department of Labour, job seekers complete a detailed daily plan for if, how and where they search. They also determine how many hours to search and applications they plan to submit per week.

Failing to follow through on the intended job search plan does not necessarily mean that people do not search at all; it could also refer to switching from high-cost, high return activities (like traveling into town to look for work or compiling and submitting applications as they risk getting rejected) to low cost activities (like browsing the job websites or calling a friend and asking whether they know of any jobs). In fact, we find that at the beginning of the study participants spend as much time searching as they intend to. By contrast, the number of applications they submit is much lower than their stated goal.

Our first main result is that five to twelve weeks after the action plan intervention, job seekers change their job search intensity, but *only* for the behavior for which we document an intention-behavior gap. Specifically, completing an action planning increases the number of job applications submitted by about 15% compared to participants who only attended the workshop, but does not affect the number of hours people spent searching. This suggests that the planning prompt increases the *efficiency* of search. That is, the action plan increases time spent on job search activities that result in submitted applications, while job-seekers in the control group appear to spend more hours on ineffective search that does not lead to applications. Self-reported data is further corroborated by observed data on job search behavior: after being informed about a vacancy seekers who completed an action plan are 11% and 29% more likely to submit an application than the workshop and control group, respectively (although the former result is not statistically significant).

Our second main finding is that the action plan induces people to use a wider range of search channels. In particular, participants switch from predominantly using informal channels like talking to family and friends to also including formal channels like submitting applications to advertisements which arguably requires more planning. This diversification of job channels leads to an increase in search *effectiveness*, arguably because of decreasing returns to using each channel and as some search activities are complementary.

These gains in search efficiency and search effectiveness translate into an increase in job offers (30%) and employment (26%), although only the former result is significantly different from the workshop group. We find that participants in the action plan group are significantly more likely to report in the follow-up survey that they found a job by answering advertisements, which is consistent with the explanation that a diversification from informal to formal channels increases search effectiveness.

Rogers et al., 2015 review the most prominent reasons suggested in the literature as to why plan-making prompts work. First, unpacking complex tasks into specific activities (when, where, and how) may help people to follow through on intentions, as they can anticipate challenges and have a more realistic understanding of the required steps. Second, planning can help overcome forgetfulness by promoting recall of the intended behavior. Simply asking people about their intentions to carry out a specific behavior increases the likelihood of follow-through as it brings tasks to their ''top of mind" (Karlan et al., 2016) or as people are more responsive to environmental cues (Gollwitzer, 1999). Finally, making a concrete action plan may serve as a commitment device. To the extent that individuals try to avoid the discomfort of failing to achieve a goal - either in front of peers or themselves - they are less likely to procrastinate (Laibson, 1997; Prestwich et al., 2012).

In order to explore which of these mechanisms may be more relevant in the domain of job search, our experimental design includes two sub-treatments. First, participants in the treatment group were further randomized to receive weekly SMS reminders about their stated

job search plans. While the reminder increases the likelihood that participants remember their search intentions by about 40%, this treatment does not improve the effectiveness of the action plan.

Second, half of the group that completed the action plan was also asked to nominate a person who could help the job seekers to follow on their search plans. This peer subsequently received text messages about the job seeker's search intentions. We find that peers are willing to serve this role and participants are very positive about how helpful this person was. However, this sub-treatment does not increases the effectiveness of the action plan either. While not conclusive, results from these two sub-treatments provide suggestive evidence against the importance of commitment (or accountability) and limited attention in our context.

The results in this paper point to action plans playing the role of ''unpacking" the different components of a multifaceted task. In particular, some sub-tasks may have higher psychological costs than others (e.g. browsing the internet vs. submitting applications and face rejection) and may differ in their returns.[51] For a given amount of time devoted to search, individuals may be more prone to focus on low-cost/low-return activities in the absence of a concrete plan that breaks up job search into well-specified and less costly tasks.[52] Consistent with this explanation, we find that the specific goals participants set themselves after completing a detailed plan of their job search activities are a significant predictor of the subsequent change in search behavior.

Our study makes several important contributions to the literature. First, we extends research on action planning to the important domain of job search. Research in social psychology suggests that the stronger the intentions to perform a certain behavior, the higher the likelihood

[51]For instance, preparing an application and submitting it to job vacancy advertisements may require the same amount of time as contacting close relatives to inquire about jobs in their firms, but the probability of receiving a response and ultimately an offer as a result of these efforts may differ significantly.

[52]The effect of i) committing to a high return but unpleasant task, and ii) figuring out how to overcome the logistical and scheduling constraints, may be important factors in overcoming the intention-behavior barrier.

of overlooking factors like plan-making, as people mistakenly believe that the strength of intentions will bring about engagement in the desired behavior (Koehler et al., 2011). Testing the effect of planning prompts for other important behaviors that are similarly complex as job search appears promising.

Our paper relates to the larger literature on the effectiveness of active labor market policies (ALMPs), and more specifically of interventions aimed to boost job search intensity and/or efficacy (Card et al., 2015). McKenzie, 2017 provides a recent reviews on the effectiveness of ALMPs in developing countries. He concludes that the vast majority of studies find modest employment gains of about 2 percentage points. In line with this conclusion, we estimate that the government-run job counseling workshop increases employment by 1.9 percentage points. Results in this paper demonstrate how simple design tweaks (e.g. adding an action plan module to a workshop) addressing behavioral biases of job seekers may improve the effectiveness of ALMPs (Babcock et al., 2012). Plan-making prompts are particularly promising as they are low-cost, easy to implement and preserve people's freedom of choice (Sunstein and Thaler, 2008).

We further contribute to an established literature investigating returns to different search channels (Holzer, 1988; Kroft and Pope, 2014; Kuhn and Mansour, 2014). Our experimental design at least partially addresses concerns about observational studies related to the endogeneity of search channel usage. Our results indicate that there are high returns to diversifying your search strategy, complementing recent experimental evidence by Belot et al., 2015 who find large returns to extending job search to additional sectors.

The rest of the paper proceeds as follows. Section 2 describes the research design and identification strategy. Section 3 reports the main results and Section 4 discusses potential mechanisms. Section 5 concludes.

2.2 Study Design

2.2.1 Background and Study Sample

A rising number of young people globally are not enrolled in education, employed, or searching for work (WorldBank, 2013). This problem is particularly severe in South Africa. Almost half of youths are unemployed (StatsSA, 2016) and 65% of young people were classified as discouraged in 2014. While (youth) unemployment in South Africa is considered largely a structural problem, recent research has documented significant frictions in the labor market.[53] The South African Department of Labour is trying to address these market inefficiencies through a range of employment services including job counseling and job referrals. However, public services are severely resources constrained in the context of sluggish economic growth. This study is part of a larger agenda to test innovative programs that are inexpensive and scalable.

Our sampling frame is the Employment Services South Africa (ESSA) data base comprising of more than 550,000 job seekers collected by the South African Department of Labour. We limit our study sample to South African unemployed job seekers between the ages of 18 and 34 who registered with ESSA in the previous 18 months and live within traveling distance from the urban labour centres that were part of the study.[54]

We randomly selected job seekers who meet these criteria and contact them using the phone number provided in ESSA. In the telephone call, surveyors invite job seekers to participate in an employment service study at the local labor center on a specified day. In return, they will receive a small stipend of 30 Rand (2.5 USD) that covers their travel cost. Of individuals successfully contacted, approximately 67% agreed to participate and of those who agreed,

[53] Abel et al., 2017 show that reducing information asymmetries between hiring firms and job seekers through reference letters can improve match quality.

[54] We worked with the following centres: Krugersdorp, Sandton and Soweto. These are urban townships in the Gauteng province.

63.5% came to the labor center on the specified day. While more educated job seekers are slightly more likely to be part of our sample, gender and age do not predict whether people accept the study invitation.

Our final sample consists of 1,097 unemployed youths. Table 11 provides summary statistics: the sample is relatively educated (12.1 years of education) and almost 80% previously held a job. Participants are actively looking for work and spend about 8 hours per week on job search and incur costs of 76 Rand (about 6 USD). However, their number of applications (4.4 per months) is relatively low.

Table 11: Sample Characteristics

	N	Mean	Median	SD
Age in yrs	1097	26.69	26	4.47
Female	1097	.52	1	.5
Education (years)	1096	12.12	12	1.16
HH Size (adults)	1097	2.36	2	1.91
1=moved to Joburg	1097	.3	0	.46
1=ever had job	1097	.79	1	.4
Reservation wage, baseline	1091	3162	3000	1832
Fair Wage	1097	5800	5000	3209
Nbr Employed Friends	1097	1.88	1	2.08
Hrs/week, BL	1058	11.35	8	9.87
Apps/month, BL	1087	4.36	3	5.26
Responses/month, BL	1088	.47	0	.86
Interviews/month, BL	1087	.24	0	.54
Offers/month, BL	1088	0	0	0
Transport costs (BS)	1043	76.92	45	91.73

Notes: Sample demographic and search behavior at baseline.

2.2.2 Intervention Design

We randomly assigned participants to one of four experimental conditions: (i) Control, (ii) Workshop only (Workshop), (iii) Workshop plus Action Plan (Workshop+AP), and (iv) Workshop plus Action Plan plus Peer Support (Workshop+AP+Peer). (Balance tables are provided in the Appendix.) The Workshop intervention is the standard 90-minute career-

counseling session conducted by the Department of Labour. The career counselor covers topics such as job search strategies, CV creation, interview techniques, and access to information and resources for job search.

Layered on top of the Workshop is a job-search planning intervention where participants were provided with an action plan template designed by the research team and were invited to create their personal action plan (Workshop+AP). The 402 job seekers assigned to this treatment were asked to think about the time they have available in a typical week and fill out the job search activities they plan to do on any given day of the week. (The Action Plan templates is provided in the Appendix.) Respondents were encouraged to be realistic about their job search plans and were asked to provide specific details, i.e. the how, when, and where of their proposed activities (for example, which newspaper to read, where to travel to search for work, etc.) since increased detail about the actualization of desired behavior has been shown to improve follow-through on intentions (Rogers et al., 2015).[55] After completing the action plan, participants are also asked to list weekly goals for hours spent searching, number of identified job opportunities, and number of submitted applications. Participants took the plan home and received an additional blank action plan template in case they wanted to change their job search plan.

A random subset of 206 job seekers in the action plan group receives an additional workshop module in which they are asked to identify a peer that can help them to follow up on their plans (Workshop+AP+Peer). Respondents provided the contact information of the nominated peer. Upon consent from both the participants and peers, the peer received weekly SMS messages about the respondent's job search goals.

Within each of the three treatment groups, job seekers were further randomized to receive text-message reminders about completing their job search goals before the end of the week

[55] Activities participants listed ranged across a spectrum of job search actions: preparatory activities such as document creation and certification, identification of opportunities, networking, and delivering CVs.

(Sunday). The reminder notifications differed slightly for each treatment group. The pure Workshop group received a general reminder to apply the lessons learned in the workshop to find job opportunities and apply for jobs. The Workshop+AP group and the Workshop+AP+Peer group received the same message but with specific reminders about their personal goals for job search and application specified in the respondent's Action Plan.[56] Participants in the control group did not receive reminder notification.

2.2.3 Data

We collect data on study participants through in-person and phone interviews. Baseline data is collected in a meeting at the labor centers. Once job seekers arrive, surveyors first register people and confirm that the ID of the participant is among those scheduled for that day. Next, the baseline survey is administered through an in-person interview. It takes on average 20 minutes and includes modules on demographic information, work history and current search activities. Baseline data was collected between September and December 2015.

Two rounds of follow-up data are collected, via phone, from all participants five and twelve weeks after the intervention.[57] The attrition rate in the first and second follow-up found is 5% and 15%, respectively, and does not differ by treatment group (Appendix B).

One concern with self-reported data is that respondents may want to please the surveyor and therefore misreport outcomes. This may be particularly relevant in the context of an intervention that was designed to assist job seekers. To address this concern we supplement

[56]The content of the reminder text-message intervention for each group was the following: Workshop treatment: Dear XX.Soweto Labour Centre Reminder: Apply the steps you learned in the job search workshop about finding job opportunities and applying for jobs." Workshop Plus Treatment: "Dear XX. Soweto Labour Centre Reminder: Your action plan is to search for X hours, find X job opportunities and apply for X jobs by Sunday."

[57]Garlick et al., 2015 show that data collected via phone and in-person surveys are largely indistinguishable.

the survey data with an observed measure of job search. Specifically, participants receive a text message from a number not associated with the research project that notifies them about a vacancy and ask them to submit an application if interested.[58]

We collect a copy of completed action plans and transcribe the content. Table 12 provides statistics of the collected action plans in terms of completion rates, number of days per week during which an activity was planned, the weekly goal number of hours to spend on job search activities, the weekly goal number of job opportunities to identify, and the weekly goal number of applications to submit.

Table 12: Action Plan Descriptives

	N	Mean (All)	SD (All)	Women	Men	Pvalue
Completed AP	402	.89	.32	.92	.86	.06
Activity-days	357	3.75	2.38	3.77	3.74	.91
Goal: Hours	345	8.46	5.78	8.11	8.82	.25
Goal: Opportunities	339	10.35	6.5	10.65	10.03	.38
Goal: Applications	340	7.82	4.33	8	7.63	.43

Notes: Descriptive statistics from completed action plans. Activitiy days refer to number of days with planned activities. P-values reported for test of equal means across gender.

We construct a measure of the intention-behavior gap from baseline data and information collected from the completed action plans. Specifically, we compare search intentions listed in the action plan with actual behavior at baseline. We find that respondents on average aim to submit 6.6 more job applications per week than they do in actuality (median difference of 5.5 applications), indicating the presence of an intention-behavior gap.[59] Conversely, respondents aim to spend on average 3.4 hours *less* time on job search activities per week than what they indicate at baseline, and this difference is centered around zero (see Figure

[58]The message informed participants about a vacancy in a specific sector, whenever possible in a sector in which they worked before. For those with work experience in different sectors, we randomly picked one sector. Sectoral shares were balanced by treatment status. Applications were submitted to actual vacancies after the last follow-up survey to avoid confounding employment estimates.

[59]Baseline behavior is collected by asking In a typical week, how many applications do you submit / hours do you search". Ideally, we would have collected an intention and behavior measure for exactly the same time period, but this was not feasible as the formation of the intention was part of the the treatment.

Figure 3: Intention Behavior Gap in Job Search Behavior

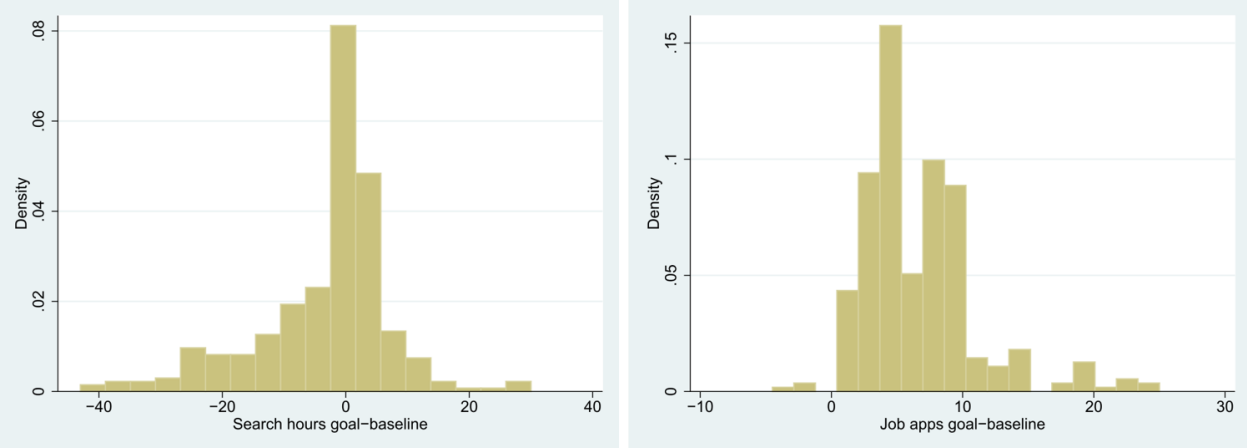

Notes: The graphs present the difference in search hours goals (as specified in the action plan) and the corresponding behavior in the last week.

3, median difference between goal hours and actual hours searched at baseline is zero).

The data thus suggests that there is an intention-behavior gap in terms of applications submitted, but not in terms of time spent searching. This will be important for the empirical analysis as it provides differential predictions about search outcomes: if the planning intervention is addressing the intention-behavior gap, then we would expect to see an effect on the number of applications and no effect on the time spent searching. By contrast, other concerns, e.g. about the self-reported nature of the data or selective attrition, would apply to both types of outcomes.

2.2.4 Empirical Strategy

The primary objective is to estimate the effects of the treatments on job search behavior and labor market outcomes. To increase statistical power, we combine the two rounds of follow-up data collection into a panel dataset and estimate the following equation for each outcome:

$$Y_{ijt} = \alpha_0 + \beta_1 Workshop_i + \beta_2 WorkshopPlus_i + \delta X_{i0} + \lambda_j + \gamma T + e_i \qquad (10)$$

where Y is the outcome indicator for individual i in location j at time t. Completion of the action plan and peer nomination was around 90%. Results presented in the next sections are intent-to-treat estimates.

As some of the characteristics were imbalanced between treatment and control groups, we report findings with and without controlling for covariate vector X. To account for geographical differences in firm demand, we control for location fixed effects j. Time dummies indicate the round of follow up (where a value of 1 signifies the second, i.e. final follow-up). Errors are clustered at the individual level to account for the panel nature of the data.

Equation 10 collapses the WS+AP and the WS+AP+Peer treatments into one treatment arm: WS Plus. Likewise, equation 10 pools across the reminder treatment. We show the effects of the Peer Support component and the Reminder treatment separately in section 2.4.

2.3 Main Results

2.3.1 Search Intensity and Efficiency

We examine job search intensity in terms of the number of hours that the respondent spends searching for a job and the number of application that the respondent completed. We do not observe a change in the number of hours spent searching in either treatment. The number of applications is significantly higher in the Workshop Plus group, and is significantly different from the pure workshop group, indicating the action planning activity is driving the effect (Table 13). An increase of 0.7 applications is equivalent to a 15% (18%) increase compared to the workshop (control) group. While this is a sizable increase in number of applications submitted relative to the low levels of baseline search activity, it only partially closes the intention-behavior gap.

Table 13: Job Search Outcomes

	Search Hours		Applications	
	(1)	(2)	(3)	(4)
WS Basic	0.225	0.016	0.163	0.124
	(0.897)	(0.897)	(0.274)	(0.273)
WS Plus	-0.243	-0.480	0.749***	0.681***
	(0.750)	(0.747)	(0.240)	(0.235)
Covariates	No	Yes	No	Yes
Observations	1888	1886	1896	1895
R^2	0.083	0.092	0.308	0.318
Control Mean	14.1	14.1	3.83	3.83
P-value	0.595	0.573	0.054	0.062

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Recall likelihood is a binary indicator whether respondents say they remember their goal. Column 2 (3) report the (absolute) difference between recalled and actual goals.

Reported increased search intensity is corroborated by measures of actual job search behavior; subjects who completed the action plan were more likely to submit an application when they are informed about a job opening. We observe a 5.3 percentage point (27%) and 2 p.p. (11%) increase in likelihood of responding compared to the workshop and control group respectively, which is qualitatively similar to the increases we observe in the self-reported number of submitted applications.

One implication of the previous results is that we see an increase in search efficiency, defined as the ratio of applications submitted to the number of hours searched. The Workshop Plus group submits 0.3 applications more per hour spent searching than the control group (a 20% increase on the control mean).

2.3.2 Employment

We investigate whether the increase in applications translate into employment related outcomes - specifically responses from employers, job offers, and employment status of the respondents. Compared to the workshop group, job seekers who completed the action plan received significantly more responses from employers regarding their application (24%), more

job offers (30%) and were more likely to be employed (26%) at the time of follow-up (Table 14).[60] These coefficients are all significantly different from both the control and workshop group with the exception of employment effects of the action plan treatment compared to the workshop group. We do not observe that the jobs that participants in the treatment group find differ with regard to salary or job satisfaction (results not reported).

Table 14: Employer Response Outcomes

	Responses		Offers		Employment	
	(1)	(2)	(3)	(4)	(4)	(5)
WS Basic	-0.025	-0.027	0.023	0.022	0.021	0.019
	(0.058)	(0.059)	(0.022)	(0.022)	(0.025)	(0.025)
WS Plus	0.112**	0.102*	0.058***	0.061***	0.047**	0.049**
	(0.053)	(0.053)	(0.020)	(0.020)	(0.020)	(0.021)
Covariates	No	Yes	No	Yes	No	Yes
Observations	1895	1894	1882	1881	1971	1969
R^2	0.101	0.109	0.012	0.021	0.017	0.024
Control Mean	0.544	0.544	0.130	0.130	0.115	0.115
P-value	0.027	0.036	0.126	0.099	0.321	0.249

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the individual level. All regressions control for location-fixed effects and baseline value of the outcome variable. Outcome variables are winsorized at the 5% level to account for outliers. P-value compares WS Plus to WS Basic.

It is noteworthy that the standard counseling program has modest positive effect of 1.9 percentage points. This result is in line with results from other programs that focus on information provision to job seekers (Altmann et al., 2015) as well as meta analyses of a broader range of of ALMPs (Card et al., 2015; McKenzie, 2017).

2.3.3 Search Strategy

It is notable that employment is moving by almost twice the margin than applications. This implies that the action plan intervention did not only lead participants to submit more

[60] Effect sizes are not significantly different by the time of the second follow-up but they are 15%-30% smaller in magnitude, suggesting that the action plan's effect may be decreasing over time. Future research could test whether a more flexible design that allows job seekers to update their action would help to facilitate sustaining the effects of action planning in the face of failed efforts.

applications, but also that each application was more effective in obtaining jobs. Applications were thus either of higher quality and/or job seekers used search channels with higher returns.

To shed light on the quality of applications, we can analyze the applications that job seekers submit in response to the vacancy notification. We find that the quality of applications seem slightly higher compared to the control group (although the difference is not significant) and very similar to the workshop group (Table 15). While participants in the action plan group seem more likely to submit formal applications, differences in the quality of application are therefore unlikely to explain the large difference in employer responses between the action plan and workshop group.

Table 15: Application Quality

	Cover Letter	Motivation	CV	Ref Letter	ID	Certificate	Quality Index
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
WS Basic	0.063	0.009	0.010	-0.026	0.043	0.019	0.119
	(0.095)	(0.053)	(0.036)	(0.019)	(0.096)	(0.094)	(0.244)
WS Plus	0.085	0.008	0.019	0.015	0.029	0.020	0.177
	(0.076)	(0.034)	(0.027)	(0.026)	(0.077)	(0.077)	(0.171)
Observations	201	201	201	201	201	201	201
R^2	0.136	0.044	0.857	0.090	0.151	0.180	0.393
Control Mean	0.365	0.068	0.878	0.014	0.378	0.392	2.095
P-value	0.800	0.989	0.776	0.075	0.874	0.994	0.807

Notes: Standard errors in parentheses.* p<0.10, ** p<0.05, *** p<0.01. Outcomes are binary variable for whether people included cover letters, motivational statements and attached various documents. The index is a sum of these six indicators.

Next we explore changes in search strategy. The Workshop Plus group diversifies its search strategy and significantly increases the use of employment agencies, dropping CVs, answering advertisements, and online search (Table 16). [61] By contrast, the workshop group does not increase the usage of any of these channels, despite the fact that the workshop covered various search strategies.

The results in Table 16 use the ordinal search frequency variable as the dependent variable, which allows the use of a simple pooled OLS estimator, but the magnitude of the estimated

[61] We find that the total number of channels in the action plan group increases by 0.24 which is significantly different from the workshop group (p-value < 0.01).

Table 16: Search Channel Use

	Empl Agency	Dropped CV	Placed Ad	Answered Ad	Searched Online	Fam/Friends
	(1)	(2)	(3)	(4)	(5)	(6)
WS Basic	-0.016	-0.140	0.060	-0.068	0.058	-0.048
	(0.139)	(0.131)	(0.144)	(0.122)	(0.130)	(0.102)
WS Plus	0.362***	0.253**	0.153	0.301***	0.410***	-0.021
	(0.123)	(0.115)	(0.119)	(0.107)	(0.100)	(0.081)
Observations	1937	1936	1934	1927	1931	1926
R^2	0.093	0.069	0.028	0.088	0.364	0.045
Control Mean	2.125	2.826	1.532	3.253	4.203	4.839
P-value	0.010	0.004	0.537	0.003	0.006	0.786

Notes: Standard errors in parentheses. Outcome variables are a categorial frequency scale from 0 (never) to 6 (daily).Regressions use panel data over two follow-up periods. Errors are clustered at the individual level.All regressions control for demographics, location-fixed effects, and baseline values of the outcome variable.Outcome variables are winsorized at the 5% level to account for outliers. P-value compares WS Plus to WS Basic.* p<0.10, ** p<0.05, *** p<0.01

coefficients do not have a simple behavioral interpretation. Appendix B addresses this shortcoming by re-estimating this model using an interval regression. The results, which can be interpreted as the effects of the interventions expressed as the additional weekly search days, are qualitatively very similar to those obtained in Table 16.

Using a diverse portfolio of search activities can be beneficial for two reasons. First, similar to other production functions, job search typically requires various inputs or activities, some of which are clearly complementary in nature. For example, identifying vacancies through online search or by visiting labor centers need to be followed by submitting applications in order to be effective. Yet, completing and submitting application material is arguably more costly - both in terms of complexity and the psychological costs of getting rejected - which may explain the substantial intention-behavior gap we document in our sample. The two search channels that do not require submitting an application, contacting family and friends, and placing advertisements, are unaffected by the Workshop Plus intervention.

Consistent with the first explanation, we find evidence suggesting that returns to spending time completing applications are high. An OLS regression using predicted search days from an interval regression finds a positive correlation between time spent answering advertisements and the number of applications submitted, the number of firm responses and employment (Table 17, Columns 1-3). In addition, the action plan group job seekers are

significantly more likely to report that they found employment by answering advertisements (Table 18). We also find generally positive effects of spending more search days on visiting employment agencies, dropping off CVs, searching online on the number of applications, firm responses and employment, whereas placing adverts or speaking to family and friends do not reveal the same positive returns. Results from a first-differenced regression (Table 17, Columns 4-6), which removes the confounding effect of time-invariant unobservable heterogeneity, are broadly consistent with the OLS results: answering adverts improves job search behavior and labor market outcomes (although the employment effect is now imprecisely estimated); the effects of dropping off CVs, searching online and visiting employment agencies remain generally (albeit less consistently) positive; and placing adverts and contacting family and friends have no positive effects.

Table 17: Effect of Predicted Search Days on Outcomes

	Applications	Responses	Employment	Applications	Responses	Employment
	(1)	(2)	(3)	(4)	(5)	(6)
Empl Agency	0.285***	0.046***	0.008*	0.100**	0.032***	-0.003
	(0.049)	(0.010)	(0.004)	(0.042)	(0.010)	(0.005)
Drop CVs	0.121**	0.022**	0.000	0.173***	0.021**	0.010**
	(0.054)	(0.011)	(0.004)	(0.045)	(0.011)	(0.005)
Place Ad	0.018	0.004	0.003	-0.004	-0.015	-0.001
	(0.051)	(0.011)	(0.004)	(0.042)	(0.010)	(0.005)
Answer Ad	0.512***	0.058***	0.008**	0.120***	0.029***	0.002
	(0.049)	(0.010)	(0.004)	(0.040)	(0.009)	(0.004)
Search Online	0.343***	0.049***	-0.001	0.077*	0.010	-0.007*
	(0.037)	(0.008)	(0.003)	(0.040)	(0.009)	(0.004)
Fam/Friends	-0.004	-0.000	-0.003	-0.013	0.004	-0.008**
	(0.037)	(0.008)	(0.003)	(0.033)	(0.008)	(0.004)
Observations	1896	1895	1922	1824	1824	882
R^2	0.197	0.092	0.007	0.025	0.019	0.014

Notes: Standard errors in parentheses. Tables report results of a regression that uses predited search days from an interval regression as explanatory variables.* p<0.10, ** p<0.05, *** p<0.01

The second reason why diversifying search activities may be effective is that each search channel has decreasing returns. Intuitively, most job seekers would start talking to the best connected friends, first search on the most effective online forum, or send the first application to the position for which they are the best fit.

Table 18: Channel used to find employment (Multinomial Regression)

	Employm. Agency	Dropped CV	Answered Advertis.	Online Search	Familie /Friends	Side of the road
	(1)	(2)	(3)	(4)	(5)	(6)
WS Basic	0.283	0.454	0.556	-0.938	-0.031	0.297
	(0.567)	(0.502)	(0.636)	(1.103)	(0.498)	(0.354)
WS Plus	0.197	0.429	1.046**	0.215	0.495	0.097
	(0.476)	(0.413)	(0.458)	(0.644)	(0.341)	(0.297)
Control Obs	Y	Y	Y	Y	Y	Y
Observations	2192	2192	2192	2192	2192	2192

Notes: Standard errors in parentheses. Tables report results of a multinomial regression. The dependent variable is an indicator for how people found their last job as reported in the follow-up survery.* p<0.10, ** p<0.05, *** p<0.01

2.4 Discussion: Why did Participants Fail to Optimize Job Search?

In the previous section, we showed that our intervention led participants to spend their job search time more efficiently and to adopt more effective search strategies. We also showed how these effects translate into substantial employment gains. This raises the question of why people failed to optimize their job search at the outset of the study. We investigate three possible barriers that the action planning may have removed: (i) complexity of a multifaceted task, (ii) forgetfulnesses, and (iii) lack of commitment. These potential mechanisms are informed by the most common explanations suggested in the planning prompts literature (Rogers et al., 2015), as briefly reviewed in the introduction. There are, of course, other mechanisms potentially at work, which renders the analysis in this section suggestive, rather than conclusive.

2.4.1 Unpacking tasks and goal setting

Action planning can help unpacking a daunting task (e.g. finding a job) into smaller sub-tasks (e.g. identify vacancies, prepare application, submit material). There are various reasons why people are more likely to follow through on smaller tasks: they require less effort, which reduces the risk of procrastination (Laibson, 1997). Also, smaller tasks tend to be more concrete, which makes individuals more likely to act in response to environmental

cues, and to have a more realistic sense of the time and logistical steps required to complete the task (Kruger and Evans, 2004; Buehler et al., 1994).

The action plan treatment asks job seekers to unpack their search process into specific actions for specific days of the week. After completing this detailed planning exercise, the job counselor asks them to set weekly goals on the number of hours spent searching, job opportunities to identify, and applications to send that are in line with their planned activities.

For participants who completed the action plan, we find a strong association between stated goals, as listed on the action plan, and subsequent behavior changes. Table 19 (Columns 2 and 3) shows that individuals who set higher goals in terms of search hours are more likely to have increased time searching for work at follow up. Specifically, a one-hour increase in the stated goal is associated with a 0.4 hour increase in actual behavior. A positive correlation is also found between goals and behavior change with respect to the number of submitted applications, with a unit increase in the stated goal being associated to a 0.11 increase in applications (Table 19, Col. 4 and 6). [62] Table 19 also shows that the intended search hours (applications) are not correlated with the subsequent change in submitting applications (hours searched). That is, goals are significant predictors of change only for the corresponding behavior, suggesting that the observed pattern is not driven by endogenous character traits (e.g. ambition).

While we cannot observe goals for individuals in the control group, these within-person correlations of job seekers in the action plan group are consistent with studies in the psychology literature showing that goal-directed actions (i.e. implementation intentions) can support goal achievement (see Brandstätter et al., 2001 for a review).[63]

[62] We do not have information on the number of identified job opportunities at follow up, so that we cannot check the correlation between goals and behavior for this variable.

[63] van Hooft and Noordzij, 2009 conduct an experiment in which they test the effectiveness of workshops that focus on goal setting. The authors find suggestive evidence of a positive effect on job intensity and employment. The study, however, is limited by a small sample size and selective attrition.

Table 19: Goal Setting

	Search Hours			Applications		
	(1)	(2)	(3)	(4)	(5)	(6)
Goal: Application	-0.131		-0.267**	0.113**		0.106*
	(0.119)		(0.119)	(0.055)		(0.054)
Goal: Hours		0.405***	0.460***		0.048	0.035
R^2	0.177	0.207	0.218	0.352	0.341	0.358
N	582	590	576	580	589	574

Notes: Standard errors in parentheses.* p<0.10, ** p<0.05, *** p<0.01.
Goals refers to the number of hours and applications job seekers specified on the action plan. Search hours and Applications refer to the actual search behavior.

2.4.2 Forgetfulness

We explore whether the effectiveness of our intervention may be explained by bringing job search to the 'top of mind' (Karlan et al., 2016). While an action plan may itself serve as a reminder, we strengthen this channel by sending weekly SMS reminders about job search goals to a subset of participants.[64]

The results reported in Table 20 show that reminders do not have significant effects on search behavior or employment outcomes. Specifically, we find that sending general reminders about job search steps to participants in the basic workshop did not increase the number of hours spent on searching or the number of applications submitted. Similarly, we find that sending specific reminders about job seekers' weekly goals (for those who had completed the action plan) does not change job search intensity within the action plan group.

One explanation for these results is that participants ignore the SMS reminders or they do not process the content of the message. However, when we call participants in the action plan group three months post-intervention, we find that reminders significantly increase the

[64]SMS reminders have been found to be effective in different domains--e.g. savings (Karlan et al., 2016) and physical activity (Prestwich et al., 2012). Other studies have used different forms of reminder to strengthen planning prompts. Milkman et al., 2013, for instance, find that merely asking people to write down the details of an medical appointment on a post-it note addresses forgetfulness leading to a 16% increases in the probability of receiving a colonoscopy.

likelihood that they could correctly recall their goal by approximately 22 percentage points (41%) (Table 21).[65]

Table 20: Reminder Effects

	(1) Search Hours	(2) Applications	(3) Empl Responses	(4) Job Offers	(5) Employed
WS Basic	-0.344	0.093	0.049	0.025	0.015
	(1.065)	(0.316)	(0.071)	(0.027)	(0.030)
WS Plus	-0.724	0.496*	0.089	0.077***	0.043*
	(0.890)	(0.282)	(0.061)	(0.026)	(0.025)
WS Basic X Reminder	0.939	0.076	-0.200**	-0.007	0.011
	(1.520)	(0.523)	(0.099)	(0.038)	(0.046)
WS Plus X Reminder	0.438	0.354	0.032	-0.031	0.012
	(0.977)	(0.366)	(0.073)	(0.030)	(0.031)
Observations	1886	1895	1894	1881	1969
R^2	0.092	0.319	0.111	0.022	0.024
Control Mean	14.095	3.835	0.544	0.130	0.115
P-value	0.741	0.271	0.603	0.100	0.391

Notes: Standard errors in parenthesesNotes: Reminder is a dummy indicator equal to 1 if observation received a reminder.Coefficient on the interaction of Reminder and treatment indicator expresses the added value of the Reminder.Regressions use panel data over two follow-up periods. Errors are clustered at the individual level.All regressions control for demographics, location-fixed effects, and baseline values of the outcome variable.Outcome variables are winsorized at the 5% level to account for outliers. P-value compares WS Plus to WS Basic.* p<0.10, **p<0.05, ***p<0.01

Table 21: Effect of Reminder on Recalling Goals

	Recall Likelihood (1)	Recall Accuracy (2)	Recall Accuracy (Abs) (3)
Reminder	0.215***	-0.151	0.037
	(0.059)	(0.876)	(0.616)
Observations	253	140	140
R^2	0.130	0.025	0.100
Control Mean	0.527	0.036	3.491

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

Recall likelihood is a binary indicator whether respondents say they remember their goal. Column 2 (3) report the (absolute) difference between recalled and actual goals.

It is also possible that the action plan itself already serves as a reminder, bringing job search to top of mind. This would limit the additional potential impact of the SMS. However, the finding that reminders are also ineffective in the group that did not complete any action plan

[65]Conditional on claiming that they remember, the actual recall of actual goal is similarly accurate across groups (Table 21, Column 3 and 4)

renders this explanation less plausible. Overall, these results suggest that forgetfulness may not be the limiting factor for the young job seekers in our sample.

2.4.3 Accountability

An action plan could serve as a commitment device as failing to follow through on intended behavior may cause discomfort (Cialdini, 2009). Receiving a reminder of your intentions could have increased accountability both towards yourself or a third part (as the reminder was addressed from the labor center), but failed to affect behavior.

Table 22: Peer Effects

	(1) Search Hours	(2) Applications	(3) Empl Responses	(4) Job Offers	(5) Employed
WS Basic	0.023	0.118	-0.028	0.022	0.020
	(0.897)	(0.273)	(0.059)	(0.022)	(0.025)
WS Plus	-0.217	0.392	0.051	0.050**	0.059**
	(0.907)	(0.305)	(0.064)	(0.025)	(0.027)
WS Plus x Peer	-0.494	0.545	0.095	0.019	-0.019
	(0.974)	(0.366)	(0.074)	(0.030)	(0.031)
Observations	1886	1895	1894	1881	1969
R^2	0.092	0.320	0.110	0.021	0.024
Control Mean	14.095	3.835	0.544	0.130	0.115
P-value	0.814	0.439	0.261	0.288	0.196

Notes: Standard errors in parenthesesNotes: Peer is a dummy indicator equal to 1 if observation is in the Peer Support treatment group.Coefficient on Peer expresses the added value of the Peer Support component.Regressions use panel data over two follow-up periods. Errors are clustered at the individual level.All regressions control for demographics, location-fixed effects, and baseline values of the outcome variable.Outcome variables are winsorized at the 5% level to account for outliers. P-value compares WS Plus to WS Basic.* p<0.10, ** p<0.05, *** p<0.01

A second test of the importance of accountability comes from the peer support intervention. We posit that the discomfort of failing to follow through on plans is greater when others are informed of the sated intentions.[66] Table 22 shows that the peer support intervention did not appear to increase job search intensity beyond the impacts of the action plan itself.

[66]Study participants mainly nominated friends (48%) and nuclear family members (42%) to receive periodic notifications about the job seeker's search goals. These close peers may be particularly instrumental in creating a sense of accountability and social pressure.

That is, participants who received an action plan had similar outcomes whether or not they nominated a peer to help them follow through on their job search intentions.

It is possible that the nominated people were not willing to serve as "support peers". When we informed peers that they were nominated by the job seeker, we provided them with a number they could text to if they did not want to fill that role. None of the support peers refused.

We also find evidence that participants engaged with their peers: 85% report in the follow-up survey that the peer helped them in their job search. Participants report of various ways other than social pressure through which peers facilitated their job search, including information provision (60%), searching on the respondent's behalf (21%), and providing financial help (14%). The fact that the peer arm does not affect any search or employment outcomes suggests that neither the commitment, nor any of the other support channels facilitates job search in our context.

It could of course be that participants were already utilizing their social network in their job search before the intervention. In fact, over 95% of respondents reported discussing job search with family and friends at baseline. However, we do not find that the peer intervention is less effective for those who already frequently discussed job search before our intervention (results not reported).

While not conclusive, the evidence presented here suggests that increasing commitment and accountability is unlikely to be the mechanism through which the action plan is effective. The apparent lack of effect through accountability may be specific to the context of our study as mass-unemployment in South Africa may reduce the stigma associated with unsuccessfully searching for jobs.

2.5 Conclusion

This study extends research on planing prompts to the domain of job search. We find that completing a detailed action plan helps South African unemployed youths to follow through on their intentions and adopt a more efficient and effective search strategy, resulting in substantial employment gains. Action planning presents a low-cost, easy to implement addition to existing job information provisions programs which have typically yielded modest results.

One open question is whether the effects of the action plan can be sustained, especially if job seekers' efforts are unsuccessful. We focus on relatively short-term behavioral changes. Future research should explore whether and how effects of action planning can be sustained. One possibility is to design online apps that allow a more flexible implementation of planning prompts.

A second question is how the social network can be leveraged to make action planning even more successful. Peer supports have been shown to be effective in facilitating various behavioral changes (see Breza, 2015 for a review). In our study, support peers did not increase the effectiveness of planning prompts. Future research may experiment with different types of peers, e.g. fellow job seekers or employed people outside the immediate network.

3 Labor Market Discrimination and Sorting: Evidence from South Africa

3.1 Introduction

It is notoriously difficult to study labor market discrimination. In recent years, a large literature has emerged in which researchers send fictitious resumes to estimate discrimination along many characteristics including race, gender, age, and sexual orientation (for recent reviews see Bertrand and Duflo, 2016; Neumark, 2016; Rich, 2014). While these correspondence studies have important advantages over observational and audit studies, they can only be used for select types of vacancies (Pager, 2007) and have several methodological limitations (Heckman, 1998).

In particular, audit studies have been applied almost exclusively in the formal sector of developed countries.[67] Yet, a minority of the global labor force works in these markets and these are precisely the settings in which employees enjoy the most comprehensive legal protection from discrimination. This study provides evidence for a very different setting: the situation of immigrants working in the informal sector in South Africa. Obtaining evidence on discrimination through audit studies is difficult in these markets as employers typically neither advertise job vacancies nor use formal hiring processes that include systematic screening of job seekers.[68] Pager, 2007, p.111 therefore concludes that in this context "*any study would require in-person application procedure*". I will argue that the observational data collected in this study allows me to not only apply many of the methodological advantages of audit studies, but also address some of the key methodological limitations.

[67] Bertrand and Duflo (2016) note, for example, that there is not a single audit study conducted in Africa.

[68] Moreover, ethical concerns have been raised about conducting audit studies that involve sending thousands of fictitious CVs in labor markets with mass unemployment as this would further impede the hiring process

I collect a unique data set of 5,500 job seeker advertisements in the domestic work, nanny and general work sector collected from South Africa's largest classified ad website, which caters mainly to the informal sector[69]. When employers search for workers in a particular sector and region they are presented with a list of truncated summaries of job seeker profiles. Based on the information visible at the search stage employers decide whether to click on the full profile, which provides additional information and contact details of the job seeker. Importantly, the profile page reports how many people previously clicked on the profile which serves as a measure of employer demand. This setting mimics the ideal experiment in that, similar to correspondence studies, researchers observe exactly the same information as the potential employer at the time of her decision whether to further screen a job seeker. I find that after controlling for other covariates, stating that one is an immigrant is associated with receiving about 10-20% fewer employer visits to the profile page than job seekers who state they are South African.

This unique data set also allows me to at least partly address previous studies' limitation that the pool of applicants from which employer can hire (in addition to the fictitious applicant) is not observed. This is an important concern as the size, qualification and racial composition of the applicant pool affect the cost of discrimination. Intuitively, the fewer and relatively less qualified applicants of the preferred race, the more costly are discriminatory hiring practices.[70] For the classified ad market I investigate, we observe the full set of available candidates presented to employers in her local labor market. What is more, the nationality composition presented to employers is quasi-random as the order is determined by the time

[69] According to the ILO, one of the defining characteristics of the informal sector is that "*Labour relations are based mostly on casual employment, kinship or personal and social relations rather than contractual arrangements with formal guarantees.*" According to the data I collect, 80% of job seekers employed through the website are agreed on the basis of a verbal agreement rather than a written contract. It is predominantely used by private households looking to hire: among job seekers who receive a response to their advertisement, 95% report being contacted by a private household.

[70] A related challenge to interpreting audit studies is that results are consistent with a world in which employers' racial preferences are proportionate to the population's racial composition but in which the share of minority applicants is disproprtionately high.

of the day that job seekers post their profile. Consistent with theoretical predictions from a tournament model in which employers discriminate against foreigners, I find that being pooled with a larger share of immigrants benefits *both* foreign and South African job seekers. Employing an estimator that soley exploits variation in the composition of the choice set is a methodological contribution that may be applied to estimate preferences in other scenarios.

Exisiting audit studies have been criticized on the ground that minority workers may be able to identify discriminating employers, which would alleviate the effects of employer discrimination (Heckman, 1998. In an extension of Becker, 1957, Arrow, 1973 shows that the wage differential between groups can disappear if there are a sufficient number of non-discriminating employers. In equilibrium, this may result is a (partially) segregated employment without (wage) discrimination. The extent to which (wage) differentials disappear depends on how well minorities can target non-discriminating employers in their job search (Black, 1995; Lang and Lehmann, 2012). On the website, I observe both where job seekers live and the location where they search for work. This data allows me to shed some light on whether minority applicants adjust their search behavior in response to discrimination. Employing a residential location fixed effect strategy, I find that immigrants search 21% further away and suggestive evidence that higher levels of discrimination faced in the area of residence induces immigrants to search for work in a different suburb. This is, to the best of my knowledge, the first evidence on geographic sorting in response to discrimination. It presents an important additional cost of labor market discrimination to job seekers, especially in a setting with high transport costs like South Africa.[71]

Last, this study hopes to contribute to the literature on the nature of discrimination.[72]

[71] Kerr, 2015 estimates that household spend on average 11% of their income on transport cost and that modes like trains or buses impose an effective tax rate of 25-30% on hourly wages.

[72] The theory of statistical discrimination was pioneered by Phelps, 1972 who reasoned that in a world of asymmetric information, employers assess the expected productivity of workers according to the average of the population with similar observable characteristics. Labor market discrimination can result when observable characteristics like race or gender are used by employers to infer information about productivity.

Altonji and Pierret, 2001, AP henceforth try to distinguish between statistical and taste discrimination by testing how the effect of observable characteristics on wage develop over time. According to the statistical discrimination theory, observable characteristics such as education and race should explain less of wage levels over time as employers learn about harder to observe determinants of productivity.[73] Building on the AP employer learning model, I find that employer demand *diverges* between South Africans and foreigners once more information is revealed on the applicants' profile. This is consistent with a model in which employers *positively* statistically discriminate against foreigners. Results from an anonymous survey I conduct with 208 domestic workers supports this conclusion: I find that immigrants have characteristics (e.g. age, education) favored by employers.

There are other differences between South African and foreign workers which may explain the results. Employers may refrain from selecting foreigners due to uncertainty of their legal status. However, I find that among foreign job seekers, those that indicate they have a work permit receive about 8% *fewer* profile clicks. The South African government suspects that some employers prefer hiring undocumented foreigners because they have less bargaining power and are thus more exploitative (DoL 2007). Alternatively, revealing that one has a work permit may also signal a higher reservation wage to employers. Results from the survey paint a subtle picture: undocumented immigrants do not have lower reservation wages or are paid lower wages than documented immigrants or South Africans. They are, however, less likely to know about or willing to utilize the CCMA, a widely used labour arbitration

---

By contrast, according to Gary Becker's model of taste-based discrimination, employers have discriminatory tastes and are thus willing to pay in order to not hire certain groups (Becker, 1957). While existing audit studies (Bertrand and Mullainathan, 2003; Oreopoulos, 2011) credibly show that discrimination exists in the screening process, they do not provide conclusive evidence on the source of discrimination (Lang and Lehmann, 2012). For example, a lower return to more credentials is at odds with both the taste and statistical discrimination approach (Bertrand and Mullainathan, 2003).

[73]The authors find little support for a model of statistical discrimination by race: while the coefficient on education falls, the negative coefficient for black race persists. One concern with this interpretation is that the discriminated group may face additional forms of discrimination such as being omitted in promotion decisions due to reasons of statistical discrimination, which may explain the persistent negative effect on race.

institution available to employees to take employers to court.

In sum, foreigners receive fewer profile clicks despite positive statistical discrimination and lower risks of being taken to court, pointing to the importance of taste discrimination. Consistent with this interpretation, I find that discrimination increases in the intimacy of the employer-worker relationship: it is highest in the nanny and lowest in the gardening sector.

One limitation of this research design is that, similar to most audit studies, we only observe the first step in the hiring process. Faced with a large number of applications they may use a simple heuristic in screening applicants and stop reading once they see a certain signal, in our case foreign nationality (Bertrand and Mullainathan, 2003). A review by Riach and J, 2002 concludes around 90% of discrimination occurs at this first selection stage. Kuhn and Shen, 2013 develop a screening cost model and show that narrow search strategies that pre-screen applicants are more commonly used in low-skilled sectors with large applicant pools. A coarse job screening strategy is thus highly relevant for the high unemployment, low-skill sectors context of this study (Bartos et al., 2016).[74] A second limitation is that while we observe the same information as employers, we may fail to notice subtle differences between profiles that differ between natives and immigrants. I try to address this by controlling for a rich set of profile characteristics including the number of words, spelling and grammar mistakes, the time and day of the posting, and a set of variables for information such as experience, age, and qualification. I also find that results are robust to controlling for job seeker characteristics *unobservable* to employers at the screening stage.

This study adds to a growing literature on online labor markets. These markets, which are predicted to increase rapidly in the future (Horton, 2010), are particularly prone to information asymmetries since, especially at an early stage in the hiring process, employers

[74]In addition, social-psychology theories of *unconscious* bias predict that discrimination is most pronounced when firms have limited information on job applicants (Arrow, 1998; Pager, 2007).

have very limited information of the job seeker (Pager, 2007). The study is most similar to Kuhn and Shen (2013, 2014) who use data from an online labor market in Xiamen, China to test how employer call back rates differ for job seekers with varying observable characteristics like age, gender, or education. This study also contibutes to a larger literature that uses data from online markets to estimate preferences. For example, Pope and Sydnor, 2011 analyze data from propsper.com, a peer-to-peer lending website in which borrowers create a loan listing with unverified personal information and pictures in order to request funding. The authors find that signals about the lenders' age, race, and gender from the advertised profile significantly affect the likelihood to receive loans and the interest rates paid by borrowers.

The paper proceeds as follows. Section 3.2 reviews the evidence on labor market discrimination and (undocumented) immigration in the South African context. Section 3.3 summarizes data collected from an online labor market for domestic workers. Section 3.4 discusses the identification strategy, reports the main results and offers model extensions that explore the role of the applicant pool composition and market thickness. Section 3.4.3 explores mechanisms and Section 3.5 tests if spatial sorting of job seekers is linked to employer preferences. Section 3.6 concludes.

3.2 Background: Immigration and labor market discrimination

Racial discrimination has played an important role in South Africa's history, both before and during *apartheid*. The post-*apartheid* regime implemented legislation to improve the economic situation of previously disadvantaged groups, most notably the Employment Equity Act (Act 55 of 1998) and the Broad-based Black Economic Empowerment (BEE) Act (Act 53 of 2003). Yet, the legacy of half a century of racial segregation and discrimination under apartheid has persistent effects even 20 years after the democratic transition. Firstly, while BEE created a small class of highly successful workers, the majority of businesses are still owned by the white minority. Secondly, racial tension and discrimination exists between

previously disadvantaged population groups, partly as the result of the apartheid's policies of racial division and the selective lifting of restrictions for Coloureds[75] and Asians in the 1980s (Seekings and Nattrass, 2008). Thirdly, it has been argued that apartheid's racist immigration laws, class division, de-sensitization to violence, and attitude of superiority towards the rest of Africa are partly to blame for South Africa's high degree of xenophobia that resulted in violent riots in 2008 and 2015 (Crush, 2008).

This paper focuses on the situation of immigrants in the South African labor market. At the end of *apartheid*, South Africa opened up to migration. The total number of documented immigrants increased by 60% from 1994 to 2003 and the share of Africans among immigrants increased from 25% to almost 50% over this time (DOL, 2007). In 2002, amidst increasing unemployment rates, the government passed the Immigration Act to facilitate easier access for South African employers to foreign skills while limiting labor market access for unskilled and semi-skilled immigrants. Employment opportunities in the formal sector became scarcer for immigrants due to "South Africans first" legislation, most notably in the mining sector. As a result, the proportion of foreign miners fell from 51% in 1997 to 38% in 2006 (DOL, 2007). Yet, rather than deterring immigration, many believe that the new migration policy led to a shift from legal to undocumented migration and employment of immigrants in the informal rather than formal sector. Undocumented immigration in years was further spurred by the political and economic crisis in neighboring countries, most notably Zimbabwe.

It is difficult to measure the number of immigrants living in South Africa. According to the 2001 census, 688,000 foreigners lived in South Africa, yet the number of irregular immigrants is likely much higher. The government estimated the number of undocumented immigrants to be around 600,000 but they cannot rule out that the true figure might be as high as three million (DOL, 2007). The government's response to the influx of undocumented im-

[75]The term '*Coloured*' refers to a heterogeneous ethnic group with ancestry from Europe, local tribes, West Africa, Mozambique and various places in Asia including India, Indonesia, and Malaysia. Coloureds tend to have lighter skin than the black population group and make up 9% of the population in South Africa and are mainly concentrated in the Western Cape (Seekings and Nattrass, 2008).

migrants was increased deportation: more than 3 million undocumented immigrants have been deported since 1990 (97% from SADC countries including 90% from Zimbabwe and Mozambique) and the annual number of cases tripled between 2002 and 2007 (Crush, 2011).

Large immigration flows combined with high unemployment rates have fostered xenophobic attitudes among South Africans as revealed in a survey of South Africans' attitudes towards migrants and refugees (Crush, 2008).[76] Yet, the South African Department of Labor (DoL) concludes that "*although a systematic survey of employers has not been conducted, case study evidence suggests that there is widespread preference for non-South African workers.*" (DOL, 2007). If true, this poses a puzzle: Do employers prefer hiring immigrants despite widespread xenophobic attitudes? What role does the legal status and exploitability of immigrants play? And how does discrimination vary between occupations? This paper hopes to shed light on these questions using data from the domestic worker, nanny, and general work sector in South Africa.

3.3 Data

I collect data on job seekers that use the website www.gumtree.co.za, South Africa's most widely used website for classified job advertisements.[77] Gumtree allows job seekers to post 'job wanted' advertisements for free in sectors ranging from housekeeping to engineering.[78]

[76] 84% of respondents believe that South Africa lets too many foreigners into the country, 74% favor deporting people that are not contributing economically, and two-thirds agree that rights to legal protection and police protection should never be granted to undocumented immigrants. Hostile attitudes are routed in beliefs that foreigners pose a criminal threat (48%), compete for jobs (37%), and bring diseases (29%). Xenophobic attitudes only differ slightly by employment status and tend to be somewhat stronger among whites across all income brackets (Crush 2010).

[77] Over all provinces, Gumtree has more than 30,000 advertisements posted by employers. Comparable numbers for the next most popular online sites and newspapers are 9000 by Job mail, 1000 by Career junction, and 14,000 by JunkMail.

[78] Gumtree also gives the option to pay for priority advertisements that are posted at the top of the search results. These advertisements are excluded from the analysis as this service is mainly used by agencies that act as brokers.

People specify the suburb where they search for employment, include a text with information about themselves, provide their home address and the position they are looking for and may upload a picture. Employers can either post 'job offered' advertisements or search for people by sector, region and suburb. Search results provide a list of matched profiles that include a truncated version of the text. By clicking on the truncated profile of a job seeker, the employer is directed to the individual profile site that contains the full text and contact information (see Appendix C). The profile page also has information on the number of people that previously visited this site which will serve as the main outcome variable of interest in the analysis of employer demand.

I collect advertisements posted by job seekers in the housekeeping, general work, and nanny sector between October 2012 and January 2013 in the Cape Town region.[79] Data was captured twice a day for all job seeker ads posted 12 to 24 hours before to alleviate potential selection problems as profiles of people successfully hired may be taken off the website.[80] I collect all posted information, the time of the posting and number of visits for more than 5,400 advertisements. I use only the text visible to employers on the *search result* site to encode characteristics provided by the job seeker including nationality, age and gender as well as qualification information such as whether the person has job experience or certificates, holds a driver's license, and is willing to live in the employer's house. I also manually review a random subset of job postings and quantify the number of spelling and grammar mistakes as well as flawed punctuation. The goal of this exercise is to codify all the profile information that may determine employer demand.

The main focus of this study is the domestic work sector, which provides employment to

[79]I focus on Cape Town as it is the region most highly on the job website. Other cities and more rural areas are unsuitable for this analysis as they have a much lower frequency of advertisement postings.

[80]Job seekers have the option of re-posting an advertisement once a month. This may potentially bias results as employers may not click on a profile that they recognize from a previous search. In practice, less than 5% of the profiles are re-posted using the identical advertisement and I exclude these from the analysis. More frequently, job seekers posts the same profile in a different suburb. This is less of a problem since employers are likely only looking for people searching in their specific suburb.

Table 23: Housekeeper characteristics by nationality

			Nationality				
	Mean	N	NoInf	SA	Malawi	Zimbabw	p-value
Nr profile clicks	6.61	3220	7.02	7.19	6.21	6.65	.001
Female	.515	3220	.322	.779	.509	.809	0
Male	.194	3220	.107	.221	.313	.092	0
Report age	.283	3220	.237	.516	.21	.474	0
Age (yr)	30.2	912	30.8	32.6	28.6	30.7	0
Report experience	.448	3220	.322	.484	.58	.385	0
Experience (yr)	3.95	527	4.87	5.42	3.47	3.61	.007
Reference	.328	3220	.296	.495	.336	.341	.013
Refer. phone nr	.024	3220	.026	.021	.022	.024	.948
Employer posts	.089	3220	.165	.063	.051	.04	.407
Workpermit	.048	3220	.03	.	.036	.101	0
Sleep in	.164	3220	.143	.105	.164	.207	.006
Able to drive	.041	3220	.028	.032	.068	.009	0
Certificate	.022	3220	.023	.011	.024	.019	.423
Virtuout traits	.412	3220	.352	.368	.422	.498	.002
Nr words	30.4	3220	30.1	33.9	30.4	30.2	0
Picture	.045	3220	.07	.063	.027	.036	.182
Punctuation	.305	3172	.297	.361	.302	.231	025
Wrong grammar	.132	3173	.163	.18	.127	.052	0
Nr Mistakes	.162	3174	.166	.131	.147	.144	0.971
Capitalized	.104	3174	.078	.098	.121	.085	.222
N			1138	102	1422	715	

Note: The table reports mean profile characteristics of job seeker profiles compared across nationalities. P-values are reported of a test of equal means across nationality groups.

about one million people or 7% of the (employed) labor force in South Africa (Dinkelman and Ranchod, 2012). Data from the South African Labor Force Survey (LFS) confirms that domestic workers are predominately drawn from lower socio-economic classes: their average level of education (7.5 years) is lower than that of other unemployed women (8.7 years) and paid workers (11.0 years). Domestic workers tend to be older (41.5 years vs. 35.1 years of population in labor force) and more likely to be black (92% vs. national average of 77%). It is the single most important sector for women; in fact, 18% of all women with paid jobs are employed as domestic workers. A study of domestic workers in Johannesburg shows that almost half of all domestic workers are either internal or cross-border migrants (Dinat and

Peberdy, 2007).

Table 23 provides summary statistics for the gumtree domestic work sample divided by the main nationality groups. As the website does not systematically collect information of job seekers, all data reported is derived from the profile text visible in the search result. For example, the age mean is calculated from the 28.3% of Gumtree job seekers that provide information on their age. There is thus the possibility of reporting bias - an issue I will discuss below.[81]

Figure 4: Distribution of Profile Clicks

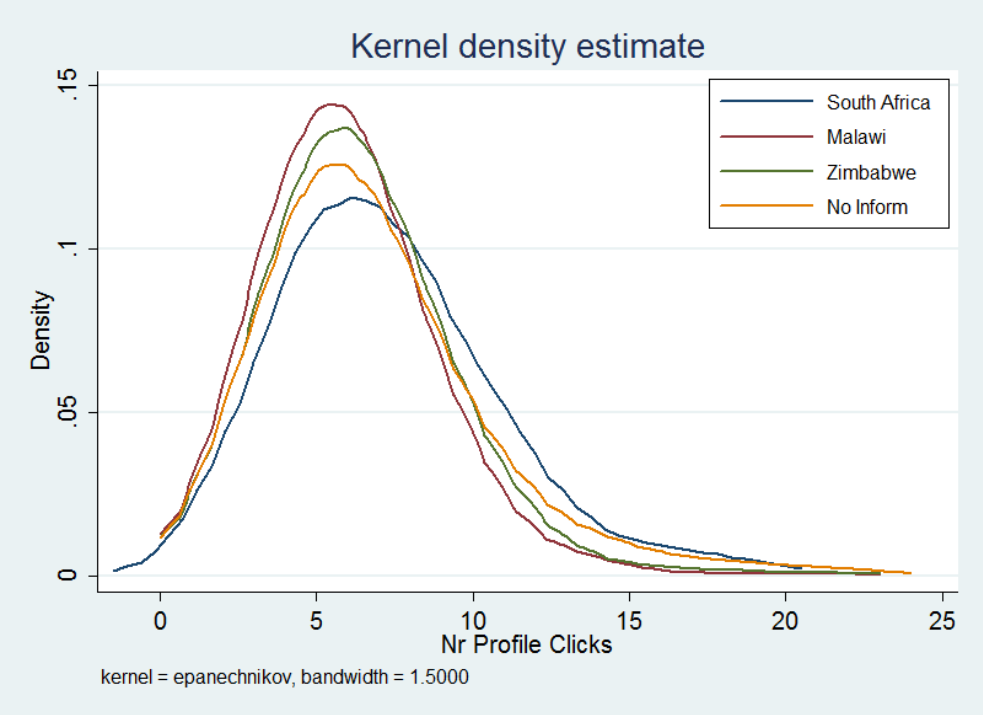

For a first descriptive summary of employer demand across population groups, Figure 3.3 plots the distribution of profile clicks for South Africans, Malawians, Zimbabweans and

[81]While this study focuses on domestic workers, I will also compare results to the nanny and general work sector (which includes gardening) to test how employer preferences vary between sectors. Table C.1 compares summary statistics across sectors. 75% and 80% of job seekers revealing their gender in the housekeeper and nanny sector, respectively, are women, whereas the general work sector is dominated by men. The share of Malawians is disproportionately high in the general work sector and Zimbabweans are overrepresented in the nanny sector. While job seeker characteristics such as age, experience, and the share holding work permits are relatively similar across sectors, I can reject a test of equal means for almost all variables.

people who do not state their nationality. The figure shows that the profile clicks follow a normal distribution for each population group. South Africans get more profile clicks than profiles without nationality information which in turn get more clicks than profiles of Zimbabweans and Malawians. The order of employer preference seems robust to outliers; in fact, the ranking of distributions (almost) follows a pattern of statistical dominace. However, these differences may be due to differences in profile characteristics between population groups documented in Table 23. The next section offers two empirical strategies to estimate the causal effect of nationality on employer demand.

3.4 Identification and Results

3.4.1 Selection on Observables

Domestic Work Sector As discussed in section 3.3, I observe exactly the same information as the employer at the time she makes the decision whether to screen an applicant's profile. While employment studies using observational data often suffer from the omitted variable problem, this study design guarantees that the observable data is orthogonal to the error term (assuming that I correctly codify the profile information). The first empirical strategy is thus to estimate what information revealed in the search result determines the decision of the employer to click on the full profile page. I estimate the following specification:

$$y_i = \alpha + \beta_1 SA + \beta_2 For + \gamma X_i + \phi_j + time_t + lang_i + e_i \tag{11}$$

I regress the log number of profile clicks y_i on a vector of control variables (X) which includes covariates listed in Table 23 such as gender, age, experience, and references. I also control for a set of suburb dummies ϕ_j and a vector of variables $time_t$ that control for the day and time at which the profile was posted and the number of hours for which the profile was online by

the time of data collection. The covariate vector $lang_i$ controls for the number of spelling, grammar and punctuation mistakes in the profile.

This paper's main focus is on the role of nationality. I include dummies capturing if the applicants are South African (SA) or a foreigner (For). I will also estimate specifications with a set of nationality dummies to account for the fact that employers may have differential preferences for applicants from different African countries. The omitted category in these specifications is the group of people providing no information on nationality. The parameters of interest β_i thus capture the difference in profile clicks among profiles with the same observable characteristics X, posted at the same time for the same suburb but with varying nationality.

Column 1 in Table 24 shows that profiles of South Africans get about 7.5% more clicks compared to profiles without nationality information whereas profiles of Zimbabweans and Malawians get 3.4% and 10.9% fewer clicks, respectively. The difference in coefficients for natives and foreigners are statistically significant. Once I control for language mistakes, the coefficient on the Zimbabwe dummy becomes more negative (Column 2). This reflects that Zimbabweans tend to make fewer language mistakes. Estimates are unchanged when I control for a set of suburb dummies which suggests that differences in employer demand are not due to differences in where people search for jobs (Column 3). In Column 4, I control for the full set of control variables. The nationality coefficients stay qualitatively unchanged and coefficients on most of the control variables have the expected sign which is reassuring. For example, male job seekers receive fewer profile clicks and people with more years of experience or who are willing to live with employers receive more clicks. I estimate the effect of age non-parametrically by including dummies for four age groups which roughly correspond to age quartiles. I find a monotonic decline of clicks in age indicating that employers prefer young domestic workers. Overall, it is reassuring that nationality coefficients remain very

Table 24: Housekeeper Analysis

	(1)	(2)	(3)	(4)	(5)
South African	0.0755	0.0306	0.0195	0.0353	0.0327
	(0.0476)	(0.0395)	(0.0398)	(0.0410)	(0.0410)
Zimbabwian	-0.0337	-0.0413*	-0.0463**	-0.0409*	
	(0.0237)	(0.0214)	(0.0212)	(0.0225)	
Malawian	-0.109***	-0.105***	-0.104***	-0.0846***	
	(0.0201)	(0.0185)	(0.0191)	(0.0199)	
Foreign					-0.0689***
					(0.0180)
Male				-0.0572***	-0.0639***
				(0.0216)	(0.0213)
Female				0.0176	0.0220
				(0.0174)	(0.0174)
Age: <25				0.0332	0.0386
				(0.0332)	(0.0331)
Age: 26-30				-0.0217	-0.0161
				(0.0226)	(0.0226)
Age: 31-35				-0.0222	-0.0125
				(0.0345)	(0.0342)
Age: 35-60				-0.0557	-0.0455
				(0.0384)	(0.0378)
Experience (yr)				0.0129*	0.0132*
				(0.00715)	(0.00715)
Reference				-0.0327*	-0.0325*
				(0.0171)	(0.0171)
Employer posts				0.00163	0.00187
				(0.0467)	(0.0467)
Refer. phone				0.0993***	0.100***
				(0.0320)	(0.0321)
Picture		0.446***	0.455***	0.438***	0.442***
		(0.0449)	(0.0453)	(0.0448)	(0.0445)
Profile Syntax	N	Y	Y	Y	Y
Suburb	N	N	Y	Y	Y
Control Var	N	N	N	Y	Y
R^2	0.011	0.198	0.204	0.216	0.217
N	3218	3218	3217	3217	3217
p-value: Zim=SA	0.028	0.083	0.115	0.050	
p-value: Mal=SA	0.000	0.001	0.002	0.003	
p-value: For=SA					0.012

Notes: standard errors in parentheses. Dependent variable: Log Profile Clicks

Base group for nationality are job seekers who do not reveal their nationality. *Employer posts* refers to whether a former employer posts on their behalf. *Reference phone* captures whether a phone nr of a reference is provided. *Reference* captures any other forms of references mentioned.

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

stable when controlling for demographic and socio-economic variables.[82]

Cross-sector comparison As described in section 3.3, I collect data from three sectors: domestic work (Dom), nanny (Nan) services, and general work (Gen). These professions are very different with respect to the frequency and intimacy of interaction. Domestic workers have access to the personal living space and often live with the employer. The relationship with nannies is even more intimate as they take care of the employer's children. If employers dislike being around foreigners, we would expect discrimination to be largest in the nanny sector and smallest in the general work sector. To test how employer demand for foreigners varies, I estimate equation 11 separately for each sector. Coefficients reported in Table 25 provide support for this hypothesis. The effect of being Malawian or Zimbabwean is about 50% and 200% larger in the domestic work and nanny sector, respectively, compared to the general work sector. However, this test remains inconclusive since there may be other unobservable differences like reputation of different nationality that vary across sectors.

The coefficients on adding a picture to the profile also shows interesting cross-sectoral differences. While the effect is large and highly significant in each sector, it is 30% and 40% larger for the domestic work and nanny sector, respectively, compared to the general work sector. Taken at face value, this indicates that the increased level of familiarity and trust that pictures evoke are more important in sectors with a more intimate employer-employee relationship.[83]

[82]While we focus on the immigrant status, there are additional interesting results, e.g. on references. 31% of profiles in the domestic work sector mention a reference in the search result. I divide these profiles into three groups: ads that were posted by the employer on behalf of their current worker (*Employer posts*), ads that list the phone number of a reference (*Refer.phone)* and those that only lists having a reference (*Reference*). Results in Column 4 show that just mentioning a reference is associated with 3.5% *fewer* profile clicks whereas providing a reference without a number has no effect. By contrast, job ads posted by current employers receive about 10% more clicks. These results suggest that employers tend to be suspicious and that references need to be credible in order to be effective Abel et al., 2017. In addition, many job seekers claim on their profiles to be 'hardworking', 'reliable' and 'trustworthy'. The effect of these claims on the number of clicks is very small and statistically insignificant (results not reported), possibly because these traits are not easily verifiable and may thus be regarded as 'cheap talk'.

[83]In a seperate analysis, I asked South Africans to rate these pictures according to two metrics: i)

Table 25: Cross Sector Analysis

	Full Sample	Housekeeper	Nanny	General
	(1)	(2)	(3)	(4)
South African	0.0335	0.0324	0.0688	0.00791
	(0.0337)	(0.0411)	(0.0814)	(0.0830)
Foreign	-0.0753***	-0.0670***	-0.135***	-0.0252
	(0.0141)	(0.0180)	(0.0322)	(0.0397)
Male	-0.0120	-0.0648***	0.0672	0.0415
	(0.0163)	(0.0213)	(0.0446)	(0.0344)
Female	-0.00235	0.0224	-0.0682**	-0.0126
	(0.0143)	(0.0174)	(0.0318)	(0.0650)
Age: <25	0.0408	0.0374	-0.001000	0.0470
	(0.0249)	(0.0332)	(0.0454)	(0.0745)
Age: 26-30	-0.0145	-0.0172	0.0190	-0.0384
	(0.0195)	(0.0226)	(0.0423)	(0.0595)
Age: 31-35	-0.0155	-0.00951	0.0138	-0.118
	(0.0282)	(0.0341)	(0.0556)	(0.103)
Age: 36-60	-0.0295	-0.0443	-0.0410	0.0960
	(0.0307)	(0.0376)	(0.0522)	(0.147)
Experience (yr)	0.00138	0.0131*	-0.00586	-0.0148
	(0.00545)	(0.00716)	(0.00957)	(0.0144)
Live w Employer	0.0661***	0.0848***	0.0388	0.0303
	(0.0185)	(0.0223)	(0.0351)	(0.0870)
Workpermit	-0.0714**	-0.0537	-0.0720	-0.0323
	(0.0295)	(0.0390)	(0.0615)	(0.0652)
Urgency	0.102**	0.169***	0.0919	0.0163
	(0.0501)	(0.0607)	(0.0950)	(0.112)
Picture	0.427***	0.442***	0.478***	0.341***
	(0.0258)	(0.0444)	(0.0405)	(0.0505)
R^2	0.266	0.217	0.320	0.191
N	5338	3217	1133	988
p-value: For=SA	0.002	0.014	0.019	0.740

Notes: Robust standard errors parentheses

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

Robustness Test: Controling for Unobservables While I observe the exact information as employers by the time they decide whether to click on a profile, one may be concerned

physical attractiveness and ii) to what extent the picture reflects the image of a domestic worker. The phycial attractiveness measure was a significant positive determinant of profile clicks whereas the second measure was negatively correlated. While these results warrant further investigation they suggest that some people may review job postings for reasons unrelated to hiring workers.

that I fail to codify all relevant information as job seeker texts are multi-dimensional. To address this concern, I take advantatge of the fact that I observe *more* information than the employer by the time she makes the screening decision. Specifically, I observe the residential location of 2,120 job seekers (which is revealed on the profile page). Given South Africa's history of spatial segregation, residential locations are correlated with socio-economic variables. By observing where job seekers live, I can therefore control for factors unobserved by the employer. The rationale for this robustness test is as follows: assume that the residential location (z_i) is correlated with the number of profile clicks after controlling for other covariates, i.e. it was previously subsumed in the error term ($e_i = v_i + z_i$). Instead of specification (11), I now control for the residential location (z_i) non-parametrically by including dummies for each 0.5x0.5 grid cell:

$$y_i = \alpha + \beta_1 SA + \beta_2 For + \gamma X_i + \phi_j + time_t + lang_i + \mathbf{z_i} + v_i \tag{12}$$

This unobservable variable explains an additional 20% of variation in profile clicks compared to specification 11. The identification concern was that $E(For\, e_i) \neq 0$ leading to biased estimates of β_2. Comparing how much coefficient β_2 changes as one reduces the error term by controlling for z_i provides an indication to what extent initial results suffered from omitted variable bias. Estimating (11) and (12) separately and testing for equal coefficients, I cannot reject that the coefficients are identical (p-value: 0.52, results not reported). This supports the validity of identification strategy (11), although it is of course still possible that $E(For\, v_i) \neq 0$.

3.4.2 Identification from Applicant Pool Composition

Framework One common criticism of audit studies is that they cannot provide information on the effect of employer discrimination in equilibrium (Heckman, 1998). If the share

of discriminating firms is small relative to the share of the minority, than the differences in callback rates observed in these studies may only have a muted effect in equilibrium. As observed by Arrow, 1973, it is the attitude of the *marginal* employer that may determine first order effects of discrimination. Charles and Guryan, 2008 attempt to test the effect of discrimination on the marginal by measuring regional racist sentiments with data from the General Social Survey.

In order to derive predictions on the role of the applicant pool composition in determining the cost and prevalence of discrimination, I propose a simple framework in which the employer hires the applicant with the highest expected productivity. This framework builds on tournament models developed by Rosen, 1981 and Lazear and Rosen, 1981 and is most closely related to Charles and Guryan, 2008.[84] The contribution of this study is to provide, to my knowledge, the first direct empirical test of the role of the applicant pool composition on discrimination.

Workers can be of type s (South African) or f (foreigners) and firms derive negative utility λ from employing type f. Let's assume that there is a fixed market wage $\bar{w}_s = \bar{w}_f = \bar{w}$ and that ability of all workers follows a uniform distribution with $a \subset [\underline{a}, \bar{a}]$. Within each population type, the firm can rank applicants according to ability a. I denote the ability level of the highest ranked applicant of each type as $a^{f*}$and a^{s*}. It is straightforward to see that the probability of both South Africans and foreigners to be highest ranked, $P[a_i^s > max(a_j^{s*}, a_j^{f*} - \lambda)], \forall j \neq i$ and $P[a_i^f - \lambda > max(a_j^{s*}, a_j^{f*} - \lambda)], \forall j \neq i$ respectively, increases in the share of foreigners, denoted θ_{For}, in the applicant pool.

Empirical Test This framework offers an indirect test of discrimination. If employers penalize foreigners, then *both* immigrants and South Africans would benefit from being pooled

[84] Cornell and Welch, 1996 argue that labor market discrimination may be the result of employers' ability to extract productivity signals more accurately for applicants of their own type. By contrast, results from the present study suggest that employers display a taste for hiring natives. However, the implications of the following simple framework do not depend on the nature of the preference for a certain group.

with more foreigners and fewer South Africans. I construct θ_{For} by measuring the nationality composition of the five applicants posting directly before and the five posting after a given profile in the same suburb[85] (see Figure C) and estimate the following model:

$$y_i = \alpha + \beta_1 For + \delta_1 \theta_{For} + \delta_2 For * \theta_{For} + \gamma X_i + \phi_j + time_t + e_i \tag{13}$$

The predictions are that both δ_1 and $(\delta_1 + \delta_2)$ are positive. It is a priori unclear if foreigners or natives benefit more from being pooled with foreigners, i.e. $\delta_2 \lessgtr 0$.

To discuss the identification assumption, imagine three job seekers ordered by the time they posted on the website.

$$y_{i-1} = \alpha + \beta For_{i-1} + \gamma X_{i-1} + e_{i-1}$$

$$y_i \quad = \alpha + \beta For_i \quad + \gamma X_i \quad + e_i$$

$$y_{i+1} = \alpha + \beta For_{i+1} + \gamma X_{i+1} + e_{i+1}$$

Identifying discrimination solely from the pool composition requires that the nationality status of job seekers who just posted before or after person i is orthogonal to i's error term, i.e. $E[\sum_{j \neq i} For_j e_i] = 0$. This identification is more likely to be valid than the selection on observable strategy which requires $E[For_i e_i] = 0$ as the neighboring profiles are determined by who posted a profile minutes earlier or later.[86] Regressing the share of foreigners θ_{For} on a host of covariates provides support for the validity of the identification assumption: of the 21 covariates collected, only one (ability to drive) is significant at the 5% level. This

[85] Computing the share θ off the previous and following five profiles is based on the fact that a total of about 10 profiles are shown on a typical computer screen. Given that only a small percentage of job seekers explicitly state they are South African and that employer preference did not differ significantly between South Africans and those not stating their nationality (Table 25), I pool these two groups for the analyses (and refer to them as SA) in this section to increase precision.

[86] While, in theory, job seekers could post their profiles at strategic times (e.g. when they see few natives), this is practically infeasible as there is a delay of about 6 hours between submitting and posting of ads.

suggests that the immigrant share displayed just above and below a profile is quasi-random.

Table 3.4.2 reports results of equation (13) estimated across the three sectors. Looking at the pooled sample, we see that job seekers indeed benefit from being in a pool with more foreigners (column 1) and that South Africans benefit more, although the difference is not statistically significant (2). Being pooled only with foreigners increases the number of profile clicks for natives by about 12.2% and for immigrants by 5.6%. While the results are generally consistent across sectors, the relative gain for natives and foreigners varies (4,6,8).

Table 26: Pool Composition

	Full Sample		Housekeeper		Nanny		General Work	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Foreign	-0.0882***	-0.0554	-0.0764***	0.0361	-0.135***	-0.191**	-0.0403	0.0450
	(0.0139)	(0.0324)	(0.0171)	(0.0555)	(0.0321)	(0.0669)	(0.0382)	(0.0652)
For. Share (θ_{For})	0.0899**	0.122**	0.0144	0.124	0.107	0.0456	0.0639	0.166
	(0.0334)	(0.0458)	(0.0472)	(0.0718)	(0.0900)	(0.117)	(0.0733)	(0.104)
For. x For. Share		-0.0636		-0.185*		0.167		-0.202
		(0.0551)		(0.0856)		(0.160)		(0.127)
Profile Syntax	Y	Y	Y	Y	Y	Y	Y	Y
Suburb	Y	Y	Y	Y	Y	Y	Y	Y
Control Var	Y	Y	Y	Y	Y	Y	Y	Y
R^2	0.265	0.265	0.211	0.212	0.316	0.316	0.182	0.184
N	5338	5338	3217	3217	1133	1133	988	988

Note: Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Foreign Share measures what percentage of the 5 postings before and after each postings were placed by foreign job seekers.

Last, it is interesting to interpret the sign of the coefficients on the foreigner main effect (β_1). This estimate measures the effect of being the only foreign job seekers that the employer observes in the local applicant pool. Across sectors this coefficient is negative and statistically significant in the nanny and housekeeper sector. This finding is inconsistent with a model of heterogeneous race preferences within the employer population. For example, if 25% of employers had preferences for immigrants, we would expect a foreign job seeker who is pooled only with natives to receive *more* profile clicks. Taken at face value, the finding provides

evidence that employer preferences with regards to nationality are not heterogeneous.

Empirical strategy 13 offers a robustness test for results from specification 12 as preference for South Africans is a necessary condition for the composition of the applicant pool to have an effect. A second necessary condition is that employers compare applicants locally, i.e. the ones presented to them at a given time on the computer screen. One implication of this assumption is that the share of foreigners presented on a previous or following search result page should matter less than the share presented on the same screen (θ_{For}^{1-5}). I test this by estimating equation 13 with the immigrant share of the applicants posted 6 to 10 places above and below a given profile (θ_{For}^{6-10}). The coefficient drops from 0.089 to 0.024 and is not statistically significant (not reported, p-value: 0.525).

3.4.3 Mechanisms

What may explain the preference for South African job seekers? There are at least three channels: i) employers risk paying a fine if they are caught hiring an immigrant who is undocumented, ii) employers may receive disutility from interacting with foreigners, and iii) employers may believe that South Africans are more productive with respect to unobservables (Appendix C offers a simple theoretical framework). To shed light on these questions I administered an anonymous phone survey with 208 domestic workers in Cape Town who posted a profile on gumtree (for more information on survey procedures, see Appendix D).

3.4.4 Legal Status and Bargaining Power

Employing an immigrant entails an expected cost as employers get fined if the employment contract is monitored and the immigrant does not have a valid work permit.[87] The data

[87] As per the Labour Act (Section 49), "*anyone who knowingly employs an illegal foreigner or a foreigner in violation of this Act shall be guilty of an offence and liable on conviction to a fine or to imprisonment not exceeding one year.*"

allows to explicitly test for the role of the legal status as about 8% of immigrants state in the search result that they have a work permit. One straightforward prediction of this mechanism is that within the group of immigrants, those with work permits receive more profile clicks. I test this hypothesis by including a dummy for workpermit (*permit*) in specification (11).

Column 1 in Table 25 shows that among foreigners having a work permit in significantly *negatively* correlated with the number of profile clicks in the pooled sample. Stating that one posesses a work permit reduces the number of profile clicks by about 7 percent.[88] While these findings are surprising at first glance, they are consistent with a theory that employers prefer job seekers that are exploitable. The DoL (2007) speculates about the reasons for anecdotal evidence showing lower unemployment among immigrants: "*The usual advantages of irregular employment (low wages, vulnerability, exploitative conditions) may be at the core of this preference* [for non-South African workers]".

Consistent with this exploitation hypothesis is the finding that job seekers stating that they are willing to live with employers or expressing that they need a job urgently receive more profile clicks (Table 25). An alternative plausible theory is that employers infer from reading that somebody holds a work permit that the person demands a higher wage and reading that somebody is 'urgently' looking for work may signal that she is less likely to shirk on the job. However, results from the anonymous phone survey do not support this explanation. Undocumented immigrants do not have lower reservation wages are not willing to work longer hours than documented immigrants (Table 27).

Results in Table 27 point to an alternative explanation for the preference of undocumented over documented immigrants: bargaining power. South Africa instituted the Commission for Conciliation, Mediation and Arbitration (CCMA), an administrative tribunal at which employers can bring cases against employer mistreatment at no charge. If disputes are not

[88]This finding also provides evidence against the explanation that employers prefer South Africans because they are more likely to remain in the Western Cape area since, controlling for other factors, immigrants with a work permit can be expected to be more established and thus stay in South Africa for longer.

resolved and no arbitration is reached, cases are referred to labour courts. About 70% of the formal workforce fall under the juristriction of the CCMA. The CCMA is very visible in the media and widely used with about 120,000 cases per year, of which about 70% are about unfair dismissal (Bhorat et al., 2009). Undocumented immigrants are significantly less likely to know about the CCMA (62.8% vs. 78.9%) and less willing to take their employer to court (57.7% vs. 70.1%) than documented immigrants.[89]

Table 27: Domestic Worker Phone Survey Results, Cape Town

	N	Sample Means				p-values			
		Pooled	SA	Doc	Undoc	S=F	S=D	S=U.	D=U
Age (yrs.)	206	32.1	38.2	31.1	29.6	0	0	0	.098
Education (yrs.)	201	11.06	10.46	11.45	11.01	.013	.001	.091	.071
Max hours willing to work	202	40.9	41.5	41.2	40.4	.588	.897	.531	.526
Wage, daily (ZAR)	191	181.3	179.4	181.2	182.3	.704	.816	.694	.867
Wage, hour (ZAR, imputed)	189	23.11	23.97	22.57	23.15	.479	.336	.567	.636
Reservation wage, daily	200	185.9	185.6	185.5	186.5	.975	.981	.903	.878
Contract	197	.228	.333	.213	.188	.084	.173	.091	.691
Treated well: 0=Never,3=Always	197	2.59	2.71	2.58	2.57	.257	.25	.338	.834
How not treated well?									
.... lower pay than agreed	208	.038	0	.052	.047	.004	.044	.044	.887
.....employer rude	208	.077	.043	.091	.082	.329	.292	.364	.848
.....had to work more hours	208	.048	.022	.078	.035	.286	.138	.648	.248
Can ask for time off	192	.958	.949	.96	.962	.741	.79	.76	.961
Paid overtime	198	.47	.476	.432	.5	.924	.653	.804	.401
Heard about CCMA	197	.756	.93	.789	.628	0	.023	0	.027
Take employer to CCMA	176	.705	.947	.701	.577	0	0	0	.131
Know about minwage law	197	.33	.395	.355	.269	.285	.669	.167	.252
Should earn minwage	168	.792	.846	.8	.75	.277	.55	.232	.5
Ever negotiated wage	196	.464	.524	.453	.443	.376	.469	.402	.899
N			46	77	85				

Notes: The table reports means of responses from a phone surveys, divided by legal and immigrant status. P-values report tests of equal means between South Africans (S), Foreigners (F), Documented (D) and Undocumented (U).

Interestingly, this lower bargaining power does not translate into lower actual daily or hourly wages. It is also notable that undocumented and documented immigrants do not report

[89]It is unclear whether undocumented immigrants have access to the CCMA. The case of Discovery Health Ltd v CCMA & others sided with an immigrant whose work permit renewal process was pending.

to be treated differently by employers. And while immigrants are more likely than South Africans to report incidents of lower than agreed payments and rude behavior from employers they report positive overall treatment by employers (2.5 on a 0-3 scale), low prevalences of maltreatment (5-10%), and feel comfortable to ask for time off if sick (96%).

However differences in bargaining power cannot explain the overall preference for South Africans.They are significantly more likely to know about the CCMA (93%) and take their employer to court (95%), are more likely to hold a have a written contract (33%) and have negotiated wages (52.4%, not significant). The next section will explore whether the results can be explained by different *expectations* about the productivity of native domestic workers.

3.4.5 Statistical vs. Taste

To shed light on the nature of discrimination, I build on insights by Farber and Gibbons, 1996 and Altonji and Pierret, 2001: if firms use nationality as a proxy for productivity, then the importance of nationality should decrease as other predictors of productivity become available and are factored into employer beliefs. By contrast, if employers have the same productivity expectation of South Africans and foreigners but prefer hiring locals due to *taste*-based discrimination, we would expect the revelation of additional information to have no effect on the foreign-national gap in hiring decisions. (For a formal employer learning model see Appendix C.)[90]

I first predict the number of clicks ($\hat{y}_i$) in the pooled sample using covariate vector X_i. The residualized numbrt of profile clicks ($y_i - \hat{y}_i$) is regressed on nationality dummies interacted with a measure of the number of relevant information (I) that the search result includes[91]:

[90]Similar tests of the nature of discrimination have been conducted through audit studies, e.g. by Baert and De Pauw, 2014. Neumark, 2016 points out that it is unclear whether the information content that is varied is relevant to employers.

[91]All variables are categorized to be relevant that have a an absolute t-value>1 in specification (11).

$$y_i - \hat{y}_i = \alpha + \beta_1 SA + \beta_2 For + \delta_1 I * SA + \delta_2 I * For + \rho I + e_i \quad (14)$$

The intuition for this test is as follows: imagine there are two binary characteristics (X_1 and X_2) that equally determine productivity. Individual A has a positive signal for X_1 and a negative for X_2 while person B does not provide any information on X_1 and X_2. While both may have the same predicted productivity, person A's profile is more informative. If employers statistically discriminate we would expect employer demand to converge as more information (I) becomes available. Given that the previous analysis found that foreigners receive fewer clicks than South Africans, we would thus expect δ_2 to be positive and δ_1 to be negative. Table 28 shows results of specification (14) estimated for the pooled sample and each sector separately.

Results are consistent across most specifications. The interaction term of nationality and information (I) is positive for South Africans and negative for Malawians and Zimbabweans (columns 2, 4, 6, 8). While providing more information is beneficial for job seekers without nationality information and especially South Africans, immigrants do not benefit (sum ρ and δ_i): at the average level of profile information, job seekers without nationality information and South Africans receive an additional 0.57 (7.7%) and 1.35 (18.2%) clicks, respectively, while demand for immigrants is unaffected. The p-values reported show that these differences in coefficients (δ_1 and δ_2) are statistically significant at the 5% level in all but the general work specification. These estimates suggest that the number of profile clicks is *diverging* between South Africans and foreigners as more information becomes available to employers. Results are consistent with a model in which employers believe that foreigners are on average more productive than natives. However, in the aggregate this *positive* statistical discrimination is

Relevant information includes gender, age, experience, references, pictures, workpermit, drivers license, and whether a person is willing to sleep in. The average profile has information on 2.9 of these variables. While applicants who don't provide information on their nationality reveal on average 2.6 characteristics, more information is revealed by South Africans (3.3), Zimbabweans (3.4) and Malawians (3.1).

Table 28: Statistical Discrimination analysis

	Full Sample		Housekeeper		Nanny		General Work	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
South African	0.558	-0.935	0.246	-2.031**	1.929	-8.333*	-0.815	-0.0466
	(0.369)	(0.946)	(0.269)	(0.748)	(1.085)	(3.941)	(0.518)	(2.292)
Zimbabwian	-0.538***	-0.0695	-0.0247	0.0429	-1.671***	-0.903	0.625	0.338
	(0.160)	(0.436)	(0.180)	(0.410)	(0.300)	(0.910)	(0.708)	(2.459)
Malawian	-0.841***	-0.302	-0.265	0.0770	-1.843***	-0.656	0.437	0.1000
	(0.141)	(0.346)	(0.154)	(0.404)	(0.361)	(1.100)	(0.377)	(0.875)
Information (I)		0.197**		0.0695		0.237		-0.0494
		(0.0710)		(0.0954)		(0.131)		(0.157)
SA x Inform (I)		0.269		0.458**		2.494*		-0.176
		(0.211)		(0.175)		(1.065)		(0.556)
Zim. x Inform (I)		-0.194		-0.0468		-0.279		0.103
		(0.107)		(0.114)		(0.216)		(0.677)
Mal. x Inform (I)		-0.210**		-0.112		-0.359		0.105
		(0.0957)		(0.118)		(0.242)		(0.250)
Profile Syntax	Y	Y	Y	Y	Y	Y	Y	Y
Suburb	Y	Y	Y	Y	Y	Y	Y	Y
Control Var	Y	Y	Y	Y	Y	Y	Y	Y
R^2	0.008	0.010	0.001	0.003	0.035	0.050	0.003	0.003
N	5340	5340	3219	3219	1133	1133	988	988
p-value: *Zim x I=SA x I*		0.044		0.002		0.012		0.806
p-value: *Mal x I=SA x I*		0.028		0.001		0.009		0.637

Notes: Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Information I measures how many relevant pieces of information were revealed in the profile.

offset by the effect of taste-based discrimination.

This conclusion is supported by previous studies concluding that immigrants are *more* employable (Crush, 2008) and by evidence from the domestic worker survey: immigrants are more educated and younger than South Africans. (Section 3.4 found that employer prefer younger job seekers.)

3.5 Supply Side Response to Discrimination: Spatial Analysis

Models of discrimination and job search can be divided into discrimination with random search and discrimination with targeted search. In random search models, the discriminated groups have lower reservation wages and accept jobs with lower match quality given that

they expect to receive fewer job offers. In models of targeted search, job seekers decide where to apply after observing firms' wage offers which provide information on the level of discrimination. The question whether job seekers adjust their search strategy in response to discriminatory practices of firms has important implications for how to interpret empirical results in the existing literature. In particular, Heckman, 1998 points out that audit studies make the critical assumption that workers employ random search strategies. If job search is costly and workers are able to direct applications to non-discriminatory firms, this should increase their reservation wage, improve match quality and thus reduce racial wage gaps in equilibrium.

Partly due to a lack of data, there is very little evidence whether discrimination induces people to direct their job search. South Africa provides a context conducive to testing these questions. One of the legacies of apartheid is that workers tend to live far away from where jobs are located.[92] The average distance of African townships from the central business districts (CBDs) of the seven largest South African cities is 28 km. Combined with relatively high costs of public transport, this results in high search and commuting costs both monetary and time-wise.[93]

While it is widely argued that spatial segregation is a barrier to finding employment, there is little evidence on where people are looking for jobs. Does high transport cost induce them to only look for jobs in their vicinity or are they willing to accept higher commuting expenses and search in areas with more employment opportunities? Are immigrants more willing to travel further for work and, if so, is this decision linked to the level of discrimination they

[92]In an attempt to claim city centers and marginalize the black and coloured population, the apartheid regime forcefully removed large parts of the urban population to homelands or township outside urban areas; the most famous cases include Sophiatown in Johannesburg and District 6 in Cape Town. Townships suffered from poor infrastructure and provided few employment opportunities as they were located far from business and industry.

[93]Recent evidence shows that people with employment spend on average R215 (7.3% of net salary) per month on transport to and from work (SALDRU, 2009). The mean amount spent by the actively searching unemployed on transport costs related to the job search was R105.75 in the previous week. 42.5% of the unemployed refrain to local job search and report not spending anything.

are facing? The online job advertisement data can shed some light on this question since job seekers have to specify the suburb where they are searching and about half provide the address or postal code. In total, I collect location data of about 3,000 job seekers in the Cape Town area.

Figure 5: Spatial Analysis of Job Search (Cape Town)

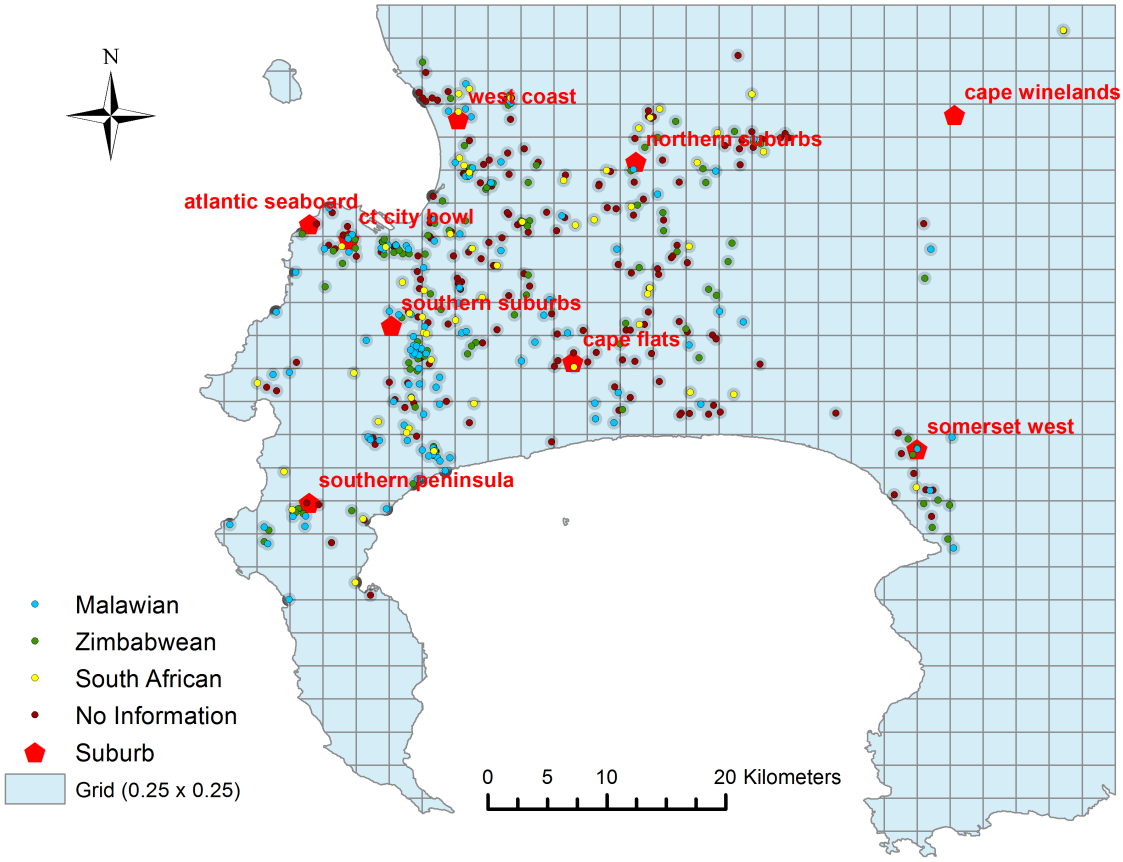

Figure 5 shows a map of the Cape Town region and the location of job seekers categorized by nationality.[94] It also displays the centroid of the nine suburbs where job seekers can indicate they are searching for employment. A few facts stand out: the distribution of job seekers roughly aligns with the population density of the urban area (with the notable exception of the poorer Cape Flats area), which lends support to the claim that the website is widely used. Second, there is no clear spatial pattern of nationality. However, one may still be

[94]To guarantee anonymity of job seekers, random noise (0.5km) is applied to the location.

concerned that difference in search behavior between immigrants and natives is linked to unobservable differences that are correlated with where people live. For example, if natives live in more central parts of the city, they would mechanically have shorter commuting distances. Conversely, if natives have better access to public transport we would expect them to search further away. These concerns are particularly important in places with large informal settlements like Cape Town where the residential location of people is linked to socio-economic factors (Hellerstein and Neumark, 2008).[95]

To address these concerns and control for unobservable differences correlated with the residential location, I draw a 0.25 x 0.25 degree gridnet (approximately 2.5 x 2.5km) and employ an empirical strategy using job seeker location fixed effects following Black et al., 2013 (Figure 5). Coefficients are estimated from variation in nationality within a grid cell with standard errors clustered at the grid cell level. First, I compare how the distance between residential location and the centroid of the suburb where job seekers look for work varies between nationalities (Panel A, Table 29). Columns 1 and 2 indicate that immigrants search for work in places that are significantly further away from their residential area: target suburbs of foreigners are on average 2.2km further compared to the pooled group of South Africans and those that do not indicate their nationality. These differences are economically meaningful given the average search distance of 10.4km. When I control for location fixed effects, coefficients decrease in magnitude but remain statistically significant. One caveat with the location fixed effect strategy is that this limits the sample used to estimate effects. I find that results are robust to extending grid cells to 0.5 x 0.5 degrees which the sample used for identification (results not reported).

[95]Previous studies pointed out that employers may discriminate based on the location where job seekers are living (Bertrand and Mullainathan 2004, Rathelot 2014). The previous results cannot be explained by this as the location of the job seeker is only revealed once employers click on their profile.

Table 29: Spatial Job Search Analysis

Panel A: **Job Search *Distance***						
	y = distance to target suburb (in m)				y=log(distance)	
	(1)	(2)	(3)	(4)	(5)	(6)
Malawian	2850.2***		1159.8**		0.0769	
	(997.2)		(485.2)		(0.0554)	
Zimbabwian	955.2		679.3*		0.103*	
	(700.7)		(383.1)		(0.0538)	
Foreigner		2243.3***		958.8***		0.0902**
		(724.7)		(343.6)		(0.0437)
Covariates	Y	Y	Y	Y	Y	Y
Location F.E.	N	N	Y	Y	Y	Y
R^2	0.043	0.037	0.440	0.440	0.658	0.658
N	2984	2984	2984	2984	2984	2984
Mean (natives)	10391	10391	10391	10391		

Panel B: **Job Search and *Home Market Discrimination***			
	y: 1=Search outside home suburb marginal effects		
	(6)	(7)	(8)
Suburb Discrimin.	0.875***	0.824***	0.831**
	(0.299)	(0.314)	(0.373)
Covariates	N	Y	Y
N	2940	2940	2940
Mean (natives)	0.281	0.281	0.281

Notes: Standard erros (clustered at the grid cell level) in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Panel B estimates standard errors through bootstrapping (200 repetitions) to account for the fact that the 'suburb discrimination' variable is estimated.

One explanation is that job seekers may decide to search further away if they are faced with discrimination in their suburb. Location subgroup analysis of specification (11) suggests that levels of discrimination vary considerably by suburb - both towards immigrants in general and between immigrant groups in a given suburb. For example, in the Northern Suburbs, Malawian profiles receive 19.3% fewer clicks than those not revealing their nationality compared to 1.6% fewer clicks for Zimbabweans. Conversely, Zimbabweans get 9% fewer clicks in Somerset West compared to 2% for Malawians. This may be the result of spatial stratification within the white population and the large historic, political and cultural differences between whites from Dutch and English descent.[96] This raises the question whether

[96]English-speaking whites live predominantly in the southern suburbs while Afrikaans-speakers more

discrimination is one of the factors that determine search patterns and could explain why immigrants search in places further away.

To shed light on this question I next test whether nationality-specific discrimination faced in the suburb of residence affects search behavior. I estimate the following regression:

$$y_i = \alpha + \delta\hat{\sigma}_{j,n} + \gamma X_i + time_t + e_j \tag{15}$$

y_i is an indicator variable capturing whether a workers searches for work outside her suburb of residence. Variable $\hat{\sigma}_{j,n}$ measures the suburb and nationality specific coefficients of discrimination estimated from equation 11. It is normalized so a higher value presents more discrimination and measures how many fewer percent profiles of nationality n receives in suburb j. To account for the facct that $\hat{\sigma}_{j,n}$ is a generated regressor, I compute standard errors (clustered at the grid cell level) through bootstrapping.

Results reported in Panel B provide suggestive evidence that discrimination in the home suburb induces migrants to search in other areas. The coefficient δ is positive and significant in column 6 and 7. Column 8 reports marginal effects estimated from a probit model. To interpret the magnitude of the coefficients, it helps to keep in mind that the aggregate coefficient of σ_n is around 0.1 (Section 3.4). Discrimination faced on the home market thus increases the probability that immigrants search in different suburbs on average by about 8 percentage points which is sizable given the mean value of 28.1% among natives.

At face value, estimates provide some of the first evidence that a group adjusts job search behavior in response to discrimination faced by employers. This has important implications for how to measure the cost discrimination. In particular it would imply that estimates of

frequently live in the northern suburbs. The imaginary dividing line between northern and southern suburbs is referred to by locals as the '*boerewors curtain*' in reference to the Afrikaans word for sausage. Furthermore, spatial stratification is reflected by the fact that the most popular newspaper in the "Die Burger" in the northern suburbs and "The Cape Argus" in the southern suburbs.

discrimination observed in Section 3.4 would underestimate the total cost of discrimination since they are estimated in equilibrium after immigrants adjusted their search behavior. The additional cost of discrimination stems from the transport cost of searching and working in places further away. However, several important caveats should be kept in mind when interpreting these results. First, coefficients ($\hat{\sigma}_{j,n}$) are estimated from samples of only a few hundred job seekers. Secondly, location choices are endogeneous and native and migrant job seekers living in the same location may differ along unobservable characteristics correlated with job search behavior. Rathelot, 2013 notes that if immigrants are penalized on the housring or labor market, the marginal immigrant to participate in the labor force should be more qualified than the marginal native. However, this should lead to *lower* commuting rates among immigrants as they have an advantage over native competitors in the same market which should reduce their inclination to search in different suburbs.

Can job seekers target suburbs with lower discrimination? To shed light on this question, I look at the subset of 1,073 commuters and measure the difference in discrimination between the home and target suburb. I find that on average levels of discrimination in the home suburb are only marginally higher (difference: 0.005 or 0.07 standard deviations, not significant). Taken at face value, this suggests that people have information on discrimination in their place of residence but that this is more difficult to assess for suburbs where one is searching for jobs.

3.6 Concluding Remarks

This study uses a unique data set of job seekers using a free job advertisement website in South Africa's Western Cape province. This is to my knowledge the first study of this kind in Africa, a continent in which the rapidly increasing internet access offers new opportunities to create and facilitate markets. In particular, internet fora like gumtree offer an opportunity to match employers with job seekers in spatially segregated urban areas and reduce the reliance

of firms to hire through social networks, currently the most common form of hiring in most developed countries (Beaman and Magruder, 2012).

One may wonder why such a large share of job seekers reveal that they are immigrants on their profile page if they get penalized by employers. One explanation is that employers can easily verify job seekers' nationality (as they typically require to see the national ID card). It is therefore beneficial for immigrants to reveal their nationality upfront to avoid going through the hiring process with discriminating employers.

One limitation of the study is that we know relatively little about the identity of potential employers using the website. From the survey we know that 96% were employed by a private household. In the Western Cape, these employers are predominantely White (63%) compared to Blacks (27%) and Coloured (10%). Whether different race groups have differential preferences over hiring immigrants vs. natives should be the subject of future research. A second limitation of the data set is that I can at best estimate discrimination at the first stage of the hiring process.Yet, in a country like South Africa with excess supply of labor and employers often reporting that they receive hundreds of applications for a single job, this first screening likely plays an important role in explaining employment outcomes.

References

Aamodt, M. (2015). *Industrial organizational psychology* (8th ed.). London: Cengage Learning.

Aamodt, M., D. Bryan, and A. Whitcomb (1993). Predicting performance with letters of recommendation. *Personnel Management 22*, 81–90.

Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin, and S. Quinn (2016). Curse of Anonymity or Tyranny of Distance? *Working Paper*.

Abel, M., R. Burger, and P. Piraino (2017). The value of reference letters: Experimental Evidence from South Africa. *Working Paper*.

Agrawal, A., N. Lacetera, and E. Lyons (2013). Does Information Help or Hinder Job Appliants from Less Developed Countries in Online Markets? *NBER Working Paper 18720*.

Altmann, S., A. Falk, S. Jaeger, and F. Zimmermann (2015). Learning about job search: A field experiment with job seekers in Germany. *Working Paper*.

Altonji, J. and C. Pierret (2001). Employer learning and statistical discrimination. *The Quarterly Journal of Economics* (February 2001).

Arrow, K. (1973). The Theory of Discrimination. In *Discrimination in Labor Markets*, pp. 3–33.

Arrow, K. (1998). What has economics to say about racial discrimination? *Journal of Economic Perspectives 12*(2), 91–100.

Autor, D. and D. Scarborough (2008). Does job testing harm minority workers? Evidence from retail establishments. *The Quarterly Journal of Economics 123*(1), 219–277.

Avery, C. and M. Meyer (2011). Designing Promotion and Hiring Procedures with Biased Evaluators. *Working Paper*.

Avery, C., P. Resnick, and R. Zeckhauser (1999). The Market for Evaluations. *The American Economic Review 89*(3), 564–584.

Babcock, L., W. J. Congdon, L. F. Katz, and S. Mullainathan (2012). Notes on behavioral economics and labor market policy. *IZA Journal of Labor Policy 1*(1), 1–14.

Baert, S. and A.-S. De Pauw (2014). Is ethnic discrimination due to distate or statistics? *Economics Letters 125*(2), 1061–72.

Bartos, V., M. Bauer, J. Chytilova, and F. Matejka (2016). Attention Discrimination: Theory and Field Experiments. *The American Economic Review 106*(6), 1437–1475.

Bassi, V. and A. Nansamba (2017). Information Frictions in the Labor Market: Evidence from a Field Experiment in Uganda. *mimeo*.

Beaman, L., N. Keleher, and J. Magruder (2013). Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi. *Working Paper* (December).

Beaman, L. and J. Magruder (2012). Who gets the job referral? Evidence from a social networks experiment. *American Economic Review 102*(7).

Becker, G. (1957). *The Economics of Discrimination*. University of Chicago Press.

Belot, M., P. Kircher, and P. Muller (2015). Does Searching Broader Improve Job Prospects? A Field Experiment. In *2015 Meeting Papers*, Number 530. Society for Economic Dynamics.

Bertrand, M. and E. Duflo (2016). Field Experiments on Discrimination. *NBER Working Paper No. 22014*.

Bertrand, M. and S. Mullainathan (2003). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. (1996).

Bhorat, H., K. Pauw, and L. Mncube (2009). *Understanding the Efficiency and Effectiveness of the Dispute Resolution System in South Africa : An Analysis of CCMA Data*. Number May. Development Policy Research Unit Working Paper 137.

Black, D. A. (1995). Discrimination in an equilibrium search model. *Journal of labor Economics 13*(2), 309–334.

Black, D. A., N. Kolesnikova, S. G. Sanders, and L. J. Taylor (2013). The role of location in evaluating racial wage disparity. *IZA Journal of Labor Economics 2*(1).

Brandstätter, V., A. Lengfelder, and P. M. Gollwitzer (2001). Implementation intentions and efficient action initiation. *Journal of personality and social psychology 81*(5), 946.

Breza, E. (2015). Field experiments, social networks, and development.

Buehler, R., D. Griffin, and M. Ross (1994). Exploring the" planning fallacy": Why people underestimate their task completion times. *Journal of personality and social psychology 67*(3), 366.

Burks, S. V., B. Cowgill, M. Hoffman, and M. Housman (2015). The Value of Hiring through Employee Referrals. *Quarterly Journal of Economics 130*(2), 805–839.

Caliendo, M., D. A. Cobb-Clark, and A. Uhlendorff (2015). Locus of control and job search strategies. *Review of Economics and Statistics 97*(1), 88–103.

Card, D., J. Kluve, and A. Weber (2015). Active Labor Market Policy Evaluations: A Meta Analysis. *NBER Working Paper No. 21431*.

Cardy, R. and G. Dobbins (1986). Affect and appraisal accuracy: Liking as an integral dimension in evaluating performance. *Journal of Applied Psychology 71*(4).

Charles, K. K. and J. Guryan (2008). Prejudice and wages: an empirical assessment of Becker's The economics of discrimination. *Journal of Political Economy 116*(5), 773–809.

Cialdini, R. (2009). Influence: The power of persuasion. *New York: William Morrow*.

Clark, D. and P. Martorell (2014). The Signaling Value of a High School Diploma. *Journal of Political Economy 122*(2), 282–318.

Cornell, B. and I. Welch (1996). Culture Information and Screening Discrimination. *Journal of Political Economy 104*(3), 542–571.

Crush, J. (2008). *The perfect storm: The realities of Xenophobia in comtemporary South Africa*. Institute for Democracy in South Africa.

Crush, J. (2011). Complex Movement, confused responses: Labour migration in South Africa.

DellaVigna, S. and M. D. Paserman (2005). Job search and impatience. Technical Report 3, National Bureau of Economic Research.

Dinat, N. and S. Peberdy (2007). Restless worlds of work, health and migration: domestic workers in Johannesburg. *Development Southern Africa 24*(1), 186–203.

Dinkelman, T. and V. Ranchod (2012). Evidence on the impact of minimum wage laws in an informal sector: Domestic workers in South Africa. *Journal of Development Economics 99*(1), 27–45.

Dipboye, R., H. Fromkin, and K. Wiback (1975). Relative importance of applicant sex, attractiveness, and scholastic standing in evaluation of job applicant resumes. *Journal of Applied Psychology 60*(1).

DOL (2007). Labour Migration and South Africa: Towards a fairer deal for migrants in the South African Economy.

Fafchamps, M. and A. Moradi (2015). Referral and Job Performance : Evidence from the Ghana Colonial Army. *Economic Development and Cultural Change 63*(4), 715–751.

Falk, A., D. Huffman, and U. Sunde (2006). Self-confidence and search.

Farber, H. S. and R. Gibbons (1996). Learning and Wage Dynamics. *Quarterly Journal of Economics 111*(4), 1007–1047.

Fields, G. (2011). Labor Market Analysis for Developing Countries. *Labour Economics 18*, 16–22.

Frey, B. S. and R. Jegen (2001). Motivation crowding theory. *Journal of Economic Surveys 15*(5), 589–611.

Garlick, R., K. Orkin, and S. Quinn (2015). Call Me Maybe: Experimental Evidence on Using Mobile Phones to Survey African Microenterprises. Technical report, Working paper, Duke University, Durham, NC.

Gibbons, R. and L. Katz (1991). Layoffs and Lemons. *Journal of Labor Economics 9*, 351–380.

Gollwitzer, P. M. (1999). Implementation intentions: strong effects of simple plans. *American psychologist 54*(7), 493.

Groh, M., D. McKenzie, and T. Vishwanath (2014). Reducing Information Asymmetries in the Youth Labor Market of Jordan with Psychometrics and Skill Based Tests. *World Bank Economic Review 29*(1), 106–117.

Grossman, S. and O. Hart (1980). Disclosure Laws and Takeover Bids. *Journal of Finance 35*(2), 323–334.

Hagger, M. S. and A. Luszczynska (2014). Implementation intention and action planning interventions in health contexts: State of the research and proposals for the way forward. *Applied Psychology: Health and Well-Being 6*(1), 1–47.

Heckman, J. (1998). Detecting discrimination. *Journal of Economic Perspectives 12*(2), 101–116.

Hellerstein, J. and D. Neumark (2008). Workplace segregation in the United States: Race, ethnicity, and skill. *The Review of Economics and Statistics 9*(3), 459–477.

Holzer, H. J. (1988). Search Method Use by Unemployed Youth. *Journal of Labor Economics 6*(1), 1–20.

Horton, J. J. (2010). *Online labor markets.* Springer Berlin Heidelberg, 2010.

Ioannides, Y. M. and L. D. Loury (2004). Job Information Networks , Neighborhood Effects, and Inequality. *Journal of Economic Literature 42*(4), 1056–1093.

Kaas, L. and C. Manger (2012). Ethnic Discrimination in Germany's Labour Market: A Field Experiment. *German Economic Review 13*(1), 1–20.

Kahn, L. and F. Lange (2014). Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures. *The Review of Economic Studies 81*(4), 1575–1613.

Karlan, D., M. McConnell, S. Mullainathan, and J. Zinman (2016). Getting to the top of mind: How reminders increase saving. *Management Science*.

Kerr, A. (2015). Tax(i)ing the poor? Commuting costs in South Africa. *Working Paper*.

Koehler, D. J., R. J. White, and L. K. John (2011). Good intentions, optimistic self-predictions, and missed opportunities. Social Psychological and Personality Science.

Kroft, K., F. Lange, and M. J. Notowidigdo (2013). Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment. *Quarterly Journal of Economics 128*(3), 1123–1167.

Kroft, K. and D. Pope (2014). Does online search crowd out traditional search and improve matching efficiency? Evidence from Craigslist. *Journal of Labor Economics 32*(2), 259–303.

Kruger, J. and M. Evans (2004). If you don't want to be late, enumerate: Unpacking reduces the planning fallacy. *Journal of Experimental Social Psychology 40*(5), 586–598.

Kuhn, P. and H. Mansour (2014). Is Internet job search still ineffective? *The Economic Journal 124*(581), 1213–1233.

Kuhn, P. and K. Shen (2013). Gender discrimination in job ads: Evidence from china. *The Quarterly Journal of Economics*, 287–336.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics 112*(2), 443–478.

Lang, K. and J.-y. K. Lehmann (2012). Racial Discrimination in the Labor Market : Theory and Empirics. *Journal of Economic Literature 50*(4), 959–1006.

Lang, K. and M. Manove (2011). Education and Labor-Market Discrimination. *American Economic Review 101*(4), 1467–1496.

Lazear, E. and S. Rosen (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 841–864.

Loher, B., J. Hazer, A. Tsai, K. Tilton, and J. James (1997). Letters of reference: A process approach. *Journal of Business and Psychology 11*(3), 339–355.

Loury, L. (2006). Some Contacts Are More Equal than Others : Informal Networks, Job Tenure, and Wages. *Journal of Labor Economics 24*(2), 299–318.

Magruder, J. (2010). Intergenerational networks, unemployment, and persistent inequality in South Africa. *American Economic Journal: Applied Economics 2*(1), 62–85.

Malindi, K. (2016). Wage Effect of Labour Market Experience and Firm Tenure for Black and White South African Workers. *Working Paper*.

McCarthy, J. and R. Goffin (2001). Improving the Validity of Letters of Recommendation: An Investigation of Three Standardized Reference Forms. *Military Psychology 13*(4), 199–222.

McGee, A. and P. McGee (2015). Search, effort, and locus of control. *Effort, and Locus of Control (July 1, 2015)*.

McKenzie, D. (2017). How Effective Are Active Labor Market Policies in Developing Countries? A Critical Review of Recent Evidence. *Working Paper*.

Milgrom, P. (1981). Rational Expectations, Information Acquisition, and Competitive Bidding. *Econometrica 49*(July), 921–943.

Milkman, K. L., J. Beshears, J. J. Choi, D. Laibson, and B. C. Madrian (2011). Using implementation intentions prompts to enhance influenza vaccination rates. *Proceedings of the National Academy of Sciences 108*(26), 10415–10420.

Milkman, K. L., J. Beshears, J. J. Choi, D. Laibson, and B. C. Madrian (2013). Planning prompts as a means of increasing preventive screening rates. *Prev Med 56*(1), 92–93.

Montgomery, B. (1991). Social Networks and Labor-Market Outcomes : Toward an Economic Analysis. *American Economic Review 81*(5), 1408–1418.

Mortensen, D. and C. Pissarides (1994). Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies 61*(3), 397–415.

Neumark, D. (2016). Experimental research on labor market discrimination. *NBER Working Paper 22022*.

Nickerson, D. W. and T. Rogers (2010). Do you have a voting plan? Implementation intentions, voter turnout, and organic plan making. *Psychological Science 21*(2), 194–199.

Oreopoulos, P. (2011). Why Do Skilled Immigrants Struggle in the Labor Market? A Field Experiment with Six Thousand Resumes. *American Economic Journal: Economic Policy 3*(4), 148–171.

Pager, D. (2007). The Use of Experiments for Studies of Employment Discrimination : Directions for the Future. *The Annals of the American Academy of Political and Social Science 609*(January), 104–133.

Pallais, A. (2014). Inefficient Hiring in Entry-Level Labor Markets. *American Economic Review 104*(11), 3565–3599.

Pallais, A. and E. G. Sands (2016). Why the Referential Treatment? Evidence from Field Experiments on Referrals. *Journal of Political Economy*.

Phelps, E. (1972, feb). The Statistical Theory of Racism and Sexism. *American Economic Review 62*(4), 659–661.

Pissarides, C. (1985). Short-Run Equilibrium Dynamics of Unemployment , Vacancies , and Real Wages. *American Economic Review 75*(4), 676–690.

Pope, D. and J. Sydnor (2011). What is in a Picture? Evidence of Discrimination from Prosper. *Journal of Human Resources 46*(1), 53–92.

Prestwich, A., M. T. Conner, R. J. Lawton, J. K. Ward, K. Ayres, and R. R. C. McEachan (2012). Randomized controlled trial of collaborative implementation intentions targeting working adults' physical activity. *Health Psychology 31*(4), 486.

Prestwich, A., R. Lawton, and M. Conner (2003). The use of implementation intentions and the decision balance sheet in promoting exercise behaviour. *Psychology and Health 18*(6), 707–721.

Rathelot, R. (2013). Ethnic differentials on the labor market in the presence of asymmetric spatial sorting: Set identification and estimation. *Mimeo CREST*.

Rees, A. (1966). Information Networks in Labor Markets. *American Economic Review 56*(1), 559–566.

Riach, P. and R. J (2002). Field Experiments of Discrimination in the Market Place. *The Economic Journal 112*, 480–518.

Riach, P. and J. Rich (2004). Deceptive Field Experiments of Discrimination: Are they Ethical? *Kyklos 57*, 457–470.

Rich, J. (2014). What do field experiments of discrimination in markets tell us? A meta analysis of studies conducted since 2000. *IZA Discussion Paper No.8584*.

Rogers, T., K. L. Milkman, L. K. John, and M. I. Norton (2015). Beyond good intentions: Prompting people to make plans improves follow-through on important tasks. *Behavioral Science & Policy 1*(2), 33–41.

Rosen, S. (1981). The economics of superstars. *American Economic Review 71*(5), 845–858.

Rospabe, S. (2001). An Empirical Evaluation of Gender Discrimination in Employment, Occupation Attainment and Wage in South Africa in the late 1990s. *Working Paper*.

SALDRU (2009). National Income Dynamics Study, Wave 1.

Schoer, V. and M. Leibbrandt (2006). Determinants of job search strategies: Evidence from the Khayelitsha/Mitchell's Plain survey. *South African Journal of Economics 74*(4), 702–724.

Schoer, V., N. Rankin, and G. Roberts (2014). Accessing the first job in a slack labour market: job matching in South Africa. *Journal of International Development 26*, 1–22.

Seekings, J. and N. Nattrass (2008). *Class, race, and inequality in South Africa*. Yale University Press, 2008.

Shepherd, D. (2008). Post-Apartheid Trends in Gender Discrimination in South Africa: Analysis through Decomposition Techniques. *Working Paper*.

Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association 13*(1), 130–167.

StatsSA (2015). National and provincial labour market: Youth. Technical Report Quartile 1., Statistics South Africa.

StatsSA (2016). Quarterly Labor Force Survey 2016. University of Cape Town. Statistics South Africa.

Sunstein, C. and R. Thaler (2008). Nudge. *The politics of libertarian paternalism. New Haven*.

Terviö, M. (2009). Superstars and Mediocrities : Market Failure in the Discovery of Talent. *Review of Economic Studies 76*(2), 829–850.

Topa, G. (2011). Labor markets and referrals. In *Handbook of Social Economics*, pp. 1193–1221.

van der Berg, S. (2007). Apartheid's Enduring Legacy: Inequalities in Education. *Journal of African Economies 16*(5), 849–880.

van Hooft, E. A. J. and G. Noordzij (2009). The effects of goal orientation on job search and reemployment: A field experiment among unemployed job seekers. *Journal of Applied Psychology 94*(6), 1581.

WorldBank (2013). World development report 2013: Jobs. World Bank.

Appendix

A Appendix Chapter 1

Tables

Table A.1: Selection: Who returns Reference Letters?

Dep var: 1=return letter	(1) demogr	(2) search	(3) aptitude	(4) job spell	(5) unemp spell	(6) job termination	(7)
Education (yr)	0.01820	0.01911	0.01221	0.01484	0.01765	0.00884	0.00568
	(0.0215)	(0.0220)	(0.0232)	(0.0217)	(0.0215)	(0.0228)	(0.0248)
Age (yr)	0.01277**	0.01272**	0.01292**	0.01767***	0.01268**	0.01370***	0.01765***
	(0.0050)	(0.0050)	(0.0051)	(0.0059)	(0.0052)	(0.0051)	(0.0059)
1=Female	-0.00344	-0.00411	-0.00686	-0.01303	-0.00345	-0.00856	-0.01830
	(0.0435)	(0.0437)	(0.0436)	(0.0441)	(0.0437)	(0.0442)	(0.0447)
Nr Applications (4 weeks)		0.00161					0.00214
		(0.0075)					(0.0076)
Aptitude Score (%)			0.00070				0.00039
			(0.0012)				(0.0013)
Last job spell (yr)				-0.02174*			-0.01762
				(0.0120)			(0.0122)
Time since last job (yr)					0.00175		0.00144
					(0.0029)		(0.0028)
Job termination: contract end						0.03234	0.03983
						(0.0508)	(0.0512)
Job termination: fired						-0.08013	-0.04502
						(0.0855)	(0.0874)
Job termination: voluntary						0.08781	0.08085
						(0.0852)	(0.0875)
R^2	0.027	0.028	0.027	0.035	0.028	0.033	0.038
N	437	435	436	437	437	437	434
Dep Var mean	0.308	0.308	0.308	0.308	0.308	0.308	0.308

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

The table explores factors correlated with whether job seekers return a completed letter. *Aptitude Score* measures the average numeracy and literacy score of an aptitude test. *Last job spell* captures the number of years the job seeker stayed in her last job. The *Job termination* variable capture the reason of termination stated by employers on the reference letter.

Table A.2: Correlation

	TeamAbility	Interpersonal	WorkEthics	Reliability	Agreeability	Numeracy	Literacy	Computer	Learning
Team Ability	1.000								
Interpersonal	0.410***	1.000							
WorkEthics	**0.550*** **	0.406***	1.000						
Reliability	0.437***	0.331***	0.416***	1.000					
Agreeability	**0.523*** **	**0.499*** **	0.431***	**0.520*** **	1.000				
Numeracy	0.276***	**0.169***	0.189**	0.372***	0.348***	1.000			
Literacy	0.308***	0.348***	0.318***	0.290***	0.381***	**0.516*** **	1.000		
Computer Lit.	**0.150**	0.321***	**0.079**	**0.134**	**0.230***	0.272***	0.344***	1.000	
Learning Ability	0.255***	0.355***	0.412***	0.301***	0.416***	0.338***	0.385***	0.187***	1.000

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The table reports correlation coefficients of employer ratings, converted to numeric values (0=below average, 3=very good). The five correlation coefficients with the largest (lowest) correlation coefficients are marked blue (red).

Table A.3: Balance Test: Reference Letter vs Control Group

	Full Sample		Control		Reference Let		pvalue
	N	mean	N	mean	N	mean	
1=Female	1267	.502	566	.516	701	.491	.373
Age in yrs	1267	27.33	566	27.07	701	27.55	.042
Education (years)	1262	12.16	561	12.08	701	12.23	.395
1=married	1267	.069	566	.055	701	.081	.06
Nr of Children	1179	1.026	525	1.021	654	1.031	.878
1=moved to Johannesburg	1267	.744	566	.753	701	.738	.539
Zulu	1267	.273	566	.281	701	.267	.575
Xhosa	1267	.084	566	.083	701	.086	.871
Venda	1267	.056	566	.049	701	.061	.356
1=ever had job	1267	1	566	1	701	1	.
1=ever selfemployed	1267	.193	566	.187	701	.197	.667
Currently receiving UIF	1267	.114	566	.102	701	.124	.225
Reservation wage (ZAR/month)	1259	3381	559	3251	700	3484	.079
Fair Wage (ZAR/month)	1265	6108	565	5930	700	6251	.143
Hours search (week)	1226	14.35	544	14.13	682	14.52	.768
Interview requests (month)	1041	.671	472	.593	569	.735	.127
Plan for job search	1132	2.972	471	2.958	661	2.982	.71
Total search cost (ZAR/month)	1107	169.01	458	168.434	649	169.416	.93
Likelihood find job	1129	2.06	471	2.038	658	2.076	.421

Notes: The table reports summary statistics for the full sample as well as separately for the control and the treatment group. The last column reports p-values of a test of equal means of the control and treatment group. Results (not reported) show that we can reject joint significance of control variables in explaining treatment status (p-value: 0.72). Likelihood find job measures preceived chances to find employment in next month (0=very unlikely, 4=very likely).

Table A.4: Attrition (Experiment 2)

	Wave 1		Wave 2	
	(1)	(2)	(3)	(4)
Reference Letter	-0.010	-0.005	-0.019	-0.017
	(0.014)	(0.014)	(0.021)	(0.021)
Education (yrs)		-0.009***		-0.009***
		(0.001)		(0.002)
Age (yrs)		-0.003**		-0.005**
		(0.002)		(0.003)
1=Female		-0.006		-0.013
		(0.014)		(0.021)
Control Variables	N	Y	N	Y
R^2	0.000	0.024	0.001	0.016
N	1246	1241	1246	1241
Control Mean	0.068	0.068	0.182	0.182

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is an indicator variable for whether people attrited in wave 1 and 2 of the follow up survey.

Table A.5: Balance Test: Take-Up Experiment

	Pooled		Control	Information		Money		Inform+Money	
	N	Mean	Mean	Mean	p-value	Mean	p-value	Mean	p-value
Age in yrs	496	26.85	27.12	27.25	.813	26.28	.134	26.73	.475
1=Female	498	.506	.508	.524	.796	.483	.697	.508	.998
Married	498	.056	.047	.056	.733	.052	.862	.069	.444
Nr of Children	498	.998	.977	1.089	.372	1.026	.711	.908	.57
Education (years)	497	11.95	11.76	12	.098	11.97	.188	12.07	.031
1=Migrant	498	.795	.781	.823	.412	.802	.695	.777	.934
1=Ever self-employed	498	.205	.227	.234	.891	.198	.591	.162	.188
Currently receiving UIF	498	.143	.109	.129	.632	.164	.22	.169	.166
Reservation wage	496	3121	2949	3299	.214	3547	.091	2738	.37
Hours search (week)	487	13.8	11.98	12.94	.555	18.08	.004	12.75	.618
Total search cost (month)	455	165.1	164	173	.71	167	.904	155	.677
Likelihood find job (month)	459	2.07	2.04	2.02	.791	2.08	.73	2.14	.283

Note: The table reports summary statistics for the pooled sample, control group and three treatment groups. P-values report results of a test of equal means of the control group and respective treatment group. *Likelihood to find job* converts reports responses converted to numeric values (0=very unlikely, 4=very likely).

Table A.6: Employer Response Effects by Sector

	Admin	Callcentre	Cleaner	Driver	Retail	Security	Unskilled
				Y=Interest			
Reference Letter	0.0115	-0.0107	0.0455	-0.0069	0.0664*	-0.0152	0.0620
	(0.0299)	(0.0320)	(0.0289)	(0.0149)	(0.0335)	(0.0087)	(0.0432)
R^2	0.142	0.087	0.047	0.072	0.070	0.019	0.128
N	429	378	272	195	380	164	232
Control Mean	0.0556	0.0599	0.0385	0.0357	0.0352	0.0159	0.0227
				Y=Interview			
Reference Letter	0.0021	-0.0110	0.0428	-0.0069	0.0556**	-0.0081	0.0089
	(0.0163)	(0.0285)	(0.0278)	(0.0149)	(0.0258)	(0.0079)	(0.0216)
R^2	0.087	0.088	0.037	0.072	0.063	0.025	0.104
N	429	378	272	195	380	164	232
Control Mean	0.0154	0.0493	0.0240	0.0357	0.0176	0.0079	0.0114

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Results presented from from Specification 4 estimated separately by sector.

Table A.7: Effect of Reference Letter Content on Call Back

	Y=Interest					Y=Interview				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Reference Letter	-0.0488	-0.0342	-0.0242	-0.0152	0.0082	-0.0666	-0.0591	-0.0432	-0.0427	-0.0142
	(0.0477)	(0.0512)	(0.0532)	(0.0541)	(0.0672)	(0.0425)	(0.0456)	(0.0446)	(0.0463)	(0.0516)
Total Score	-0.0037	-0.0062	-0.0028	0.0000	-0.0049	-0.0002	-0.0027	-0.0001	0.0030	0.0002
	(0.0068)	(0.0065)	(0.0066)	(0.0068)	(0.0069)	(0.0049)	(0.0057)	(0.0049)	(0.0044)	(0.0053)
Ref Let x Total Score	0.0123	0.0241**	0.0182*	0.0107	0.0203*	0.0160*	0.0220**	0.0208**	0.0145	0.0204*
	(0.0099)	(0.0102)	(0.0101)	(0.0109)	(0.0110)	(0.0087)	(0.0099)	(0.0102)	(0.0097)	(0.0113)
Nr of Comments	0.0015				0.0010	0.0002				-0.0004
	(0.0020)				(0.0024)	(0.0014)				(0.0017)
Nr Comm. x Ref Let	0.0059*				0.0057	0.0037				0.0022
	(0.0035)				(0.0045)	(0.0035)				(0.0036)
Confidence		0.0145			0.0144		0.0116			0.0097
		(0.0138)			(0.0140)		(0.0083)			(0.0084)
Confidence x Ref Let		-0.0337			-0.0389*		-0.0160			-0.0165
		(0.0206)			(0.0211)		(0.0170)			(0.0174)
Termination: Contract				0.0207**	0.0209**				0.0052	0.0049
				(0.0095)	(0.0096)				(0.0069)	(0.0074)
Contract x Ref Let				-0.0153	-0.0225				-0.0076	-0.0131
				(0.0205)	(0.0208)				(0.0194)	(0.0184)
Termin.: Retrenchment				0.0155	0.0136				0.0265*	0.0229
				(0.0182)	(0.0188)				(0.0138)	(0.0144)
Retrench. x Ref Let				-0.0011	-0.0006				-0.0124	-0.0185
				(0.0329)	(0.0379)				(0.0256)	(0.0288)
Termin.: Fired				0.0498	0.0505				0.0709	0.0650
				(0.0808)	(0.0772)				(0.0872)	(0.0845)
Fired x Ref Let				-0.0997	-0.0908				-0.0834	-0.0881
				(0.0893)	(0.0848)				(0.0901)	(0.0892)
N	1919	1919	1919	1919	1919	1919	1919	1919	1919	1919
Control mean	0.042	0.042	0.042	0.042	0.042	0.024	0.024	0.024	0.024	0.024

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at applicant level.

Nr of Comments measures the number of comments that are provided by employers in the skill section of the template.

Table A.8: Short-run Effect of Reference Letter on Employment (5 weeks)

	Intent to Treat Effects			**Local Average Treatment Effects**		
	(1)	(2)	(3)	(4)	(5)	(6)
	Application	Interview	Employment	Application	Interview	Employment
			Panel A: POOLED			
Reference Letter	0.462	-0.037	0.014	1.153	-0.093	0.037
	(0.382)	(0.045)	(0.016)	(0.945)	(0.112)	(0.040)
R^2	0.289	0.062	0.007	0.296	0.055	0.005
N	1120	1122	1162	1120	1122	1162
Control Mean	4.683	0.365	0.076	4.683	0.365	0.076
			Panel B: FEMALE			
Reference Letter	-0.112	-0.031	-0.005	-0.312	-0.083	-0.012
	(0.542)	(0.062)	(0.022)	(1.462)	(0.164)	(0.059)
R^2	0.393	0.068	0.014	0.344	0.063	0.012
N	564	565	589	564	565	589
Control Mean	4.748	0.356	0.073	4.748	0.356	0.073
			Panel C: MALE			
Reference Letter	0.791	-0.052	0.036	2.286*	-0.124	0.086
	(0.498)	(0.065)	(0.023)	(1.235)	(0.155)	(0.055)
R^2	0.307	0.065	0.013	0.263	0.053	0.002
N	546	556	573	546	556	573
Control Mean	4.683	0.374	0.078	4.683	0.374	0.078

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Results presented in Column 1-3 are intent to treat estimates. Results in Column 4-6 are treatment on the treated estimates, using the encouragement assignment as an instrument for take-up. Panel A reports estimates from Specification 6 for the full sample. *Application* and *Interviews* measures the nr of applications submitted and job interviews in the last four weeks, respectively. The number of applications and interviews are winsorized at the 1% level to account for outliers. Employment is an indicator variable measuring if people are in paid employment or self-employed. Panel A and B estimate results separately for women and men.

Figures

Figure A.1: Reference Letter Template

Subject: Reference for _____ _____
(Name) *(Address of Firm)*
_____
(Address of Firm)

To Whom it May Concern:

My name is_____ . I am the _____ of _____ .
(Name) *(Position)* *(Firm / Business Name)*

Our firm is _____.
(Describe what firm is doing)

I have known _____ for _____. He/She has worked for our firm as a _____ for _____.
(Name) *(Time Known)* *(Position)* *(Time Worked)*

From _____ interactions I feel _____ to accurately judge his attitude and skills.
(daily/weekly/monthly) *(very confident / confident / somewhat confident)*

Attitude	Rating					Comment
Team ability: Ability to work under supervisor and in a team.	Very good	Good	Average	Below Average	Cannot rate	
Interpersonal skills: Friendliness and communication with customers/ coworkers	Very good	Good	Average	Below Average	Cannot rate	
Work Ethics: Willingness and ability to work hard.	Very good	Good	Average	Below Average	Cannot rate	
Reliability: Show up on time and not mismanage funds / equipment	Very good	Good	Average	Below Average	Cannot rate	

ADDITIONAL COMMENT on Attitude: _____
_____

Skill	Rating					Comment
Numeracy: Math skills necessary for this job.	Very good	Good	Average	Below Average	Cannot rate	
Literacy: Reading / Writing skills needed for this job.	Very good	Good	Average	Below Average	Cannot rate	
Computer literacy: Use of Windows, Word, Excel, Internet, etc.	Very good	Good	Average	Below Average	Cannot rate	
Task 1: *(Describe Task)*	Very good	Good	Average	Below Average	Cannot rate	
Task 2: *(Describe Task)*	Very good	Good	Average	Below Average	Cannot rate	

ADDITIONAL COMMENT on Skills: _____
_____

Our employment relationship ended because _____
(Reason for end of employment)

I would _____ _____.
(highly recommend / recommend / recommend with reservations) *(Name)*

If you have any questions do not hesitate to contact me via phone _____ *AND/OR* email _____.

Sincerely,

_____ _____
Signature Date

Figure A.2: Reference Letter Template - Examples

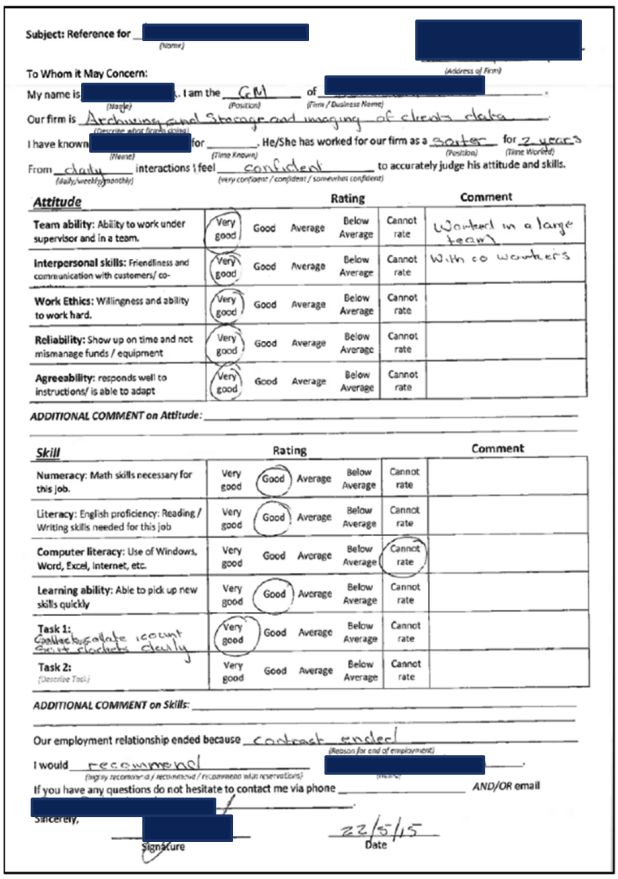

Subject: Reference for [redacted] (Name)

[redacted] (Address of Firm)

To Whom it May Concern:

My name is [redacted] (Name). I am the GM (Position) of [redacted] (Firm / Business Name).

Our firm is Archiving and Storage and imaging of clients data (Describe what firm is doing)

I have known [redacted] (Name) for _____ (Time Known). He/She has worked for our firm as a sorter (Position) for 2 years (Time Worked).

From daily (daily/weekly/monthly) interactions I feel confident (very confident / confident / somewhat confident) to accurately judge his attitude and skills.

Attitude	Rating	Comment
Team ability: Ability to work under supervisor and in a team.	(Very good) Good Average Below Average Cannot rate	Worked in a large team
Interpersonal skills: Friendliness and communication with customers/ co-	(Very good) Good Average Below Average Cannot rate	With co workers
Work Ethics: Willingness and ability to work hard.	(Very good) Good Average Below Average Cannot rate	
Reliability: Show up on time and not mismanage funds / equipment	(Very good) Good Average Below Average Cannot rate	
Agreeability: responds well to instructions/ is able to adapt	(Very good) Good Average Below Average Cannot rate	

ADDITIONAL COMMENT on Attitude: _____

Skill	Rating	Comment
Numeracy: Math skills necessary for this job.	Very good (Good) Average Below Average Cannot rate	
Literacy: English proficiency: Reading / Writing skills needed for this job	Very good (Good) Average Below Average Cannot rate	
Computer literacy: Use of Windows, Word, Excel, internet, etc.	Very good Good Average Below Average (Cannot rate)	
Learning ability: Able to pick up new skills quickly	Very good (Good) Average Below Average Cannot rate	
Task 1: Collect bank account checks daily (Describe Task)	(Very good) Good Average Below Average Cannot rate	
Task 2: (Describe Task)	Very good Good Average Below Average Cannot rate	

ADDITIONAL COMMENT on Skills: _____

Our employment relationship ended because contract ended (Reason for end of employment).

I would recommend (highly recommend / recommend / recommend with reservations) [redacted] (Name).

If you have any questions do not hesitate to contact me via phone _____ AND/OR email [redacted]

Sincerely,

[redacted] Signature

22/5/15 Date

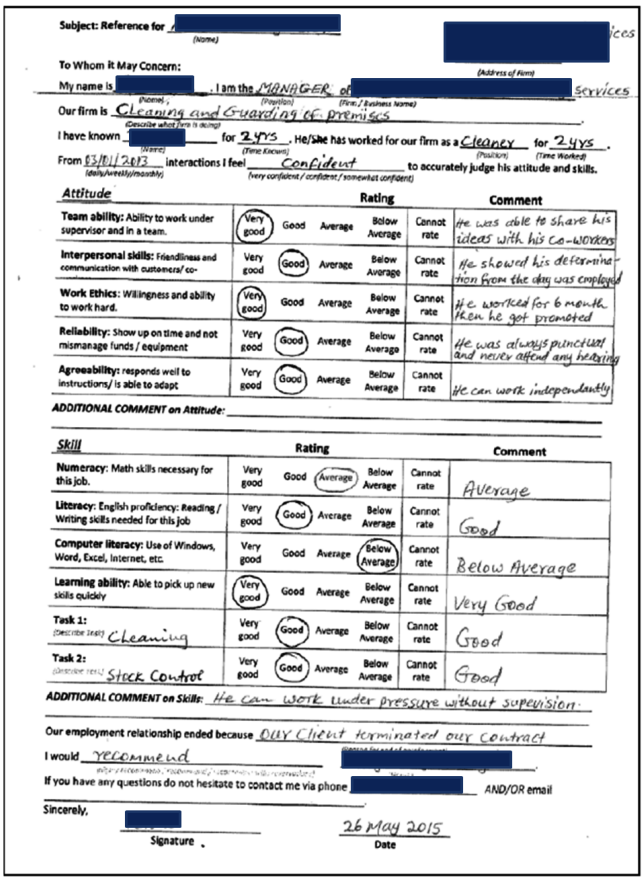

Subject: Reference for [redacted] (Name)

[redacted] (Address of Firm)

To Whom it May Concern:

My name is [redacted] (Name). I am the MANAGER (Position) of [redacted] (Firm / Business Name).

Our firm is Cleaning and Guarding of premises (Describe what firm is doing)

I have known [redacted] (Name) for 2yrs (Time Known). He/She has worked for our firm as a Cleaner (Position) for 2yrs (Time Worked).

From 03/01/2013 (daily/weekly/monthly) interactions I feel Confident (very confident / confident / somewhat confident) to accurately judge his attitude and skills.

Attitude	Rating	Comment
Team ability: Ability to work under supervisor and in a team.	(Very good) Good Average Below Average Cannot rate	He was able to share his ideas with his co-workers
Interpersonal skills: Friendliness and communication with customers/ co-	Very good (Good) Average Below Average Cannot rate	He showed his determination from the day was employed
Work Ethics: Willingness and ability to work hard.	(Very good) Good Average Below Average Cannot rate	He worked for 6 month. then he got promoted
Reliability: Show up on time and not mismanage funds / equipment	Very good (Good) Average Below Average Cannot rate	He was always punctual and never attend any hearing
Agreeability: responds well to instructions/ is able to adapt	Very good (Good) Average Below Average Cannot rate	He can work independently

ADDITIONAL COMMENT on Attitude: _____

Skill	Rating	Comment
Numeracy: Math skills necessary for this job.	Very good Good (Average) Below Average Cannot rate	Average
Literacy: English proficiency: Reading / Writing skills needed for this job	Very good (Good) Average Below Average Cannot rate	Good
Computer literacy: Use of Windows, Word, Excel, internet, etc.	Very good Good Average (Below Average) Cannot rate	Below Average
Learning ability: Able to pick up new skills quickly	(Very good) Good Average Below Average Cannot rate	Very Good
Task 1: (Describe Task) Cleaning	Very good (Good) Average Below Average Cannot rate	Good
Task 2: (Describe Task) Stock Control	Very good (Good) Average Below Average Cannot rate	Good

ADDITIONAL COMMENT on Skills: He can work under pressure without supervision.

Our employment relationship ended because Duty Client terminated our contract

I would recommend (highly recommend / recommend / recommend with reservations) [redacted] (Name).

If you have any questions do not hesitate to contact me via phone [redacted] AND/OR email

Sincerely,

[redacted] Signature

26 May 2015 Date

Figure A.3: Aptitude Test - Sample Questions

MATH

25 + 8 =

0.58 + 1.29 =

11.39 - 3.18 =

25 ÷5 =

3 +(2x 5) =

What is larger 1/4 or 1/3?

Three quarters of 100=

Which of these means 8/10: a) 80 b) 78 c) 0.8 d) 0.08

30% of R100

How many meters in a kilometre?

The time is 8:10. What time will it be in 1hours and 40 mins?

The distance to work is 50km and I am halfway there, how much longer do I still have to travel to get to work?

At the bake sale, you sold biscuits for R2 each. You earned R32. How many biscuits did you sell?

Thandeka works from 8am to 11am. Every hour, she sells 6 books. How many books does she sell in a day?

ENGLISH

Please fill in the correct word:

Alicia, _____ the windows please. It's too hot in here.

- A. opens
- B open
- C opened
- D will opened

Maria _____ never late for work.

- A. am
- B. are
- C. were
- D. is

Please read the paragraph and answer the questions:

Tomorrow, you will need to pick up the fish from the harbour. Be there in the morning before it gets too hot and the fish will go bad. You can tell the fisher that the secretary will make a bank transfer tomorrow. Please bring the fish directly to the restaurant so that the chef can use it for lunch customers.

Why do you need to pick up the fish early?

Who fill make the payment?

A) the fisher
B) you
C) the secretary
D) the chef

Where are you supposed to drop off the fish?

Figure A.4: Aptitude Distribution

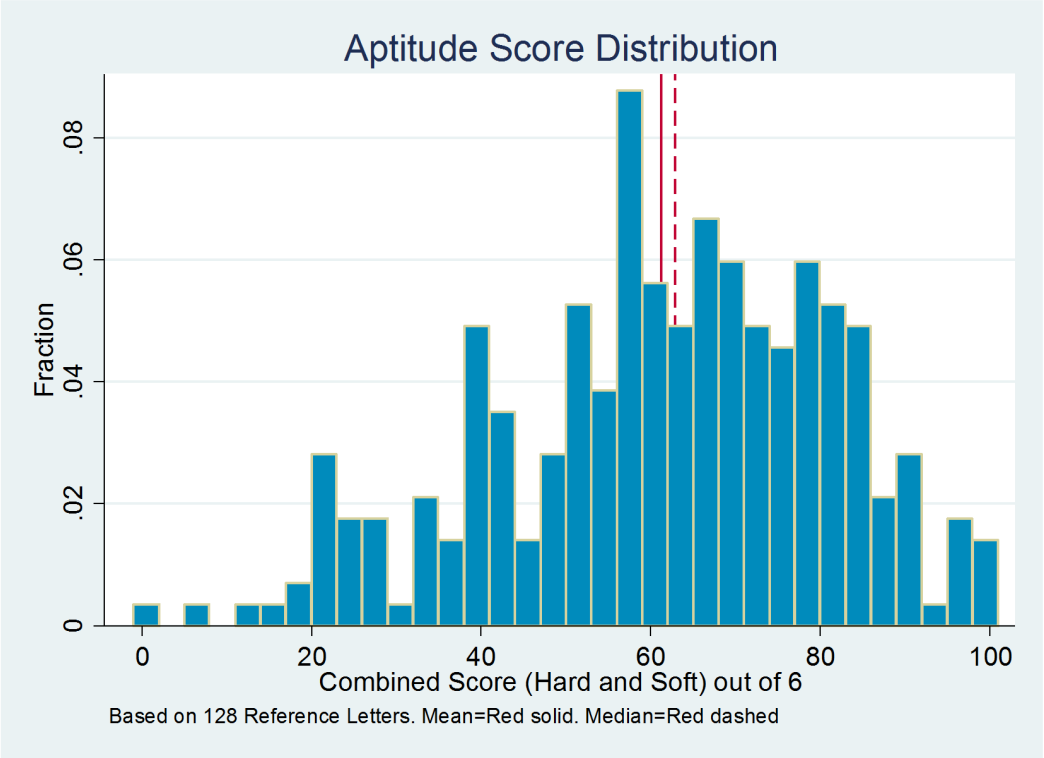

Figure A.5: Correlation between Employer Rating and Aptitude

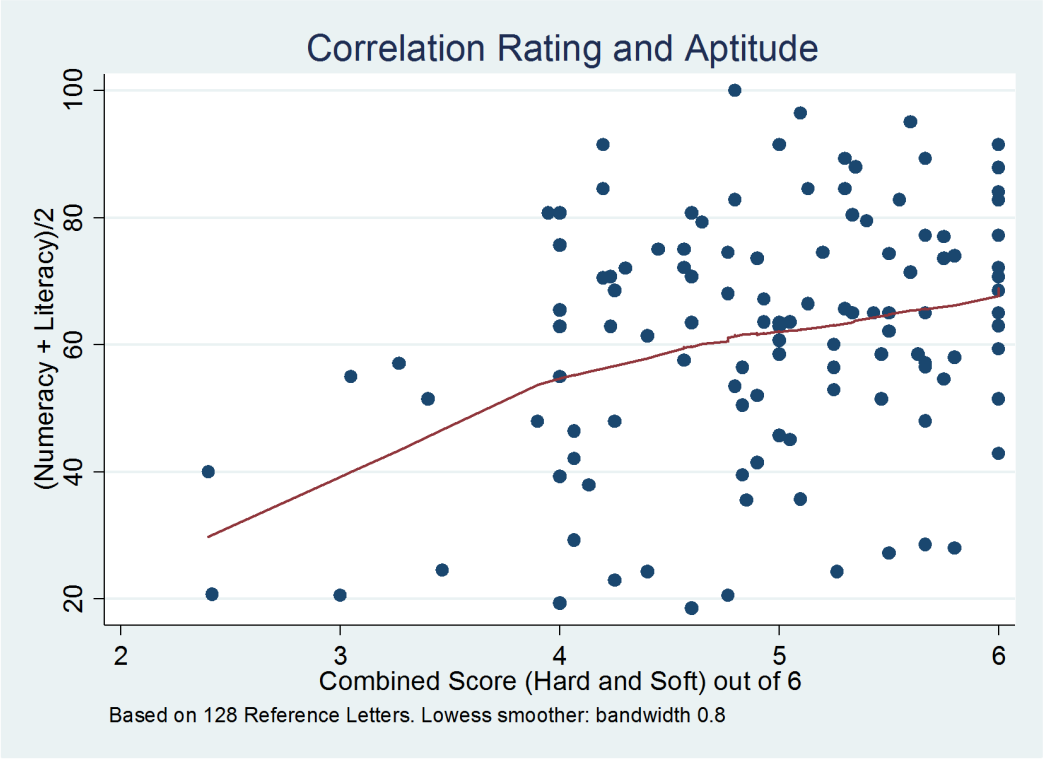

Figure A.6: Distribution of Reference Letter Scores

Figure A.7: Quartile Regression

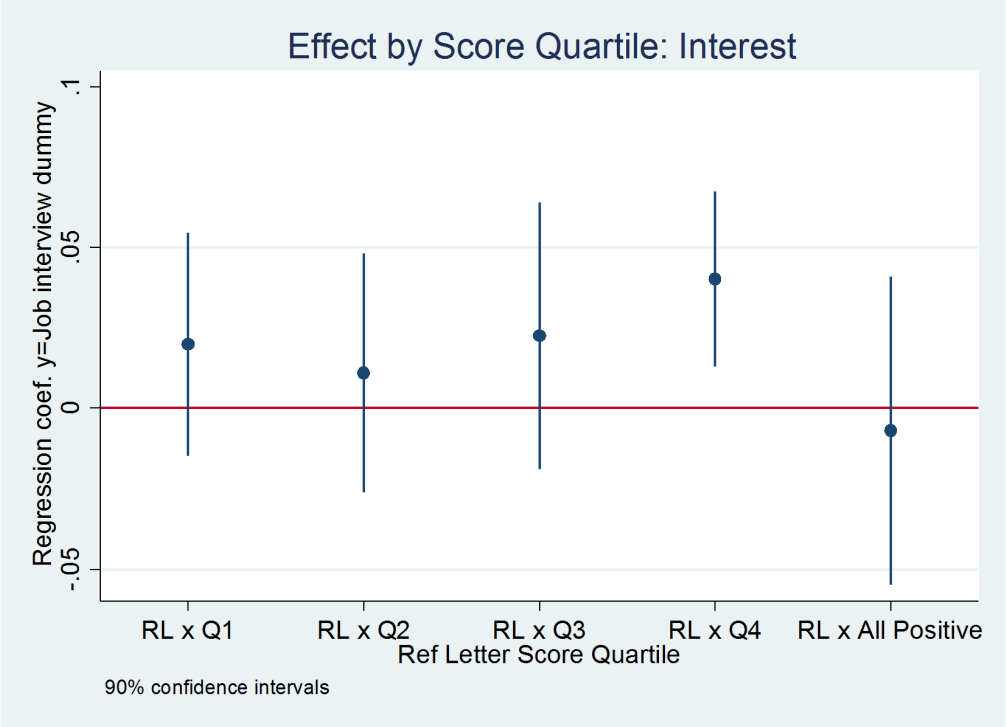

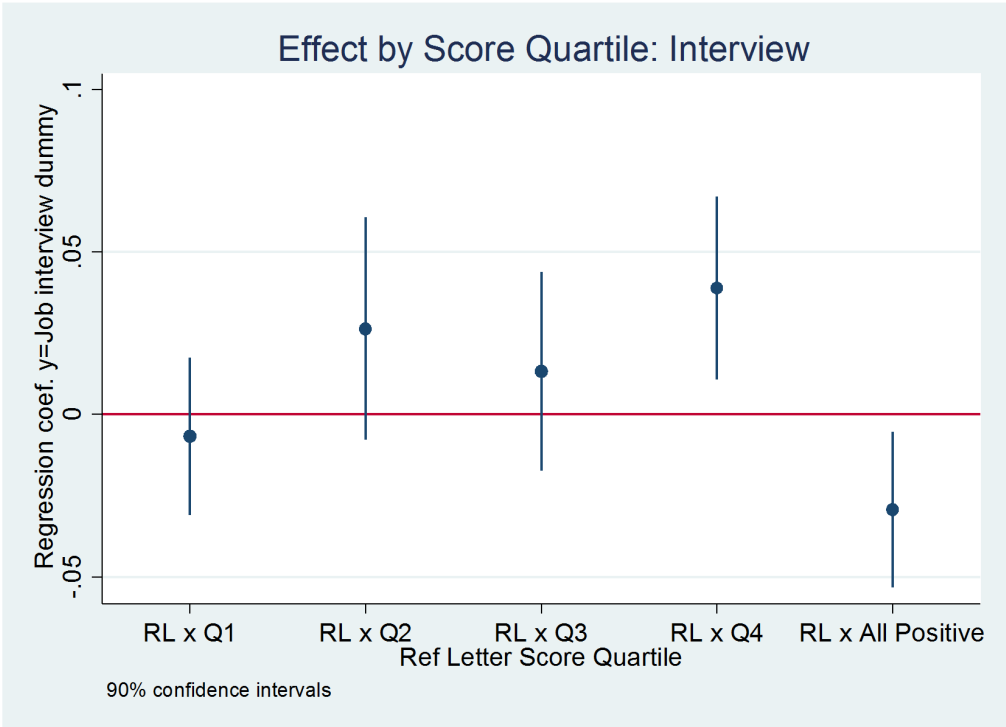

Model

The firm's conditional expectation function is

$$E(a|s_1, s_2, d) = \beta_{00} + \beta_{01}d + \beta_{10}s_1 + \beta_{11}s_1 d + \beta_{21}s_2$$

Since the firm's expectations are rational and common knowledge, the job-seeker's decision to send a reference letter is

$$d(s_1, s_2, d) = c.1(E(a|s_1, d = 0) < E(a|s_1, s_2, d = 1))$$

Under the additional linearity assumption this becomes

$$d(s_1, s_2,) = c.1(\beta_{01} + \beta_{11}s_1 + \beta_{21}s_2 > 0) = c.1(s_2 > -\frac{\beta_{01}}{\beta_{21}} - \frac{\beta_{11}}{\beta_{21}}s_1)$$

If the candidate sends a reference letter, then the employer observes s_1 and s_2 but no additional information about a is conveyed by the fact that the letter was sent. The linear coefficients of $E(a|s_1, s_2, d = 1)$ can therefore be calculated as regression coefficients. The linear regression coefficients for

$$E(a|s_1, s_2, d = 1) = (\beta_{00} + \beta_{01}) + (\beta_{10} + \beta_{11})s_1 + \beta_{21}s_2$$

can be calculated via the Frisch-Waugh-Lovell theorem as

$$\beta_{10} + \beta_{11} = \frac{\sigma_2^2}{\sigma_2^2 + \sigma_1^2\sigma_2^2 + \sigma_1^2}$$

$$\beta_{21} = \frac{\sigma_1^2}{\sigma_2^2 + \sigma_1^2\sigma_2^2 + \sigma_1^2}$$

$$\beta_{00} + \beta_{01} = 0$$

so that

$$E(a|s_1, s_2, d = 1) = \frac{\sigma_2^2}{\sigma_2^2 + \sigma_1^2\sigma_2^2 + \sigma_1^2}s_1 + \frac{\sigma_1^2}{\sigma_2^2 + \sigma_1^2\sigma_2^2 + \sigma_1^2}s_2$$

However, when no reference letter is sent the employer should use this information to update their expectation about the value of s_2. By the law of iterated conditional expectations:

$$E(a|s_1, d) = E(E(a|s_1, s_2, d)|s_1, d) = \beta_{00} + \beta_{01}d + \beta_{10}s_1 + \beta_{11}s_1d + \beta_{21}dE(s_2|s_1, d)$$

The expected value of s_2 given the observed value of s_1 and the fact that no reference letter was sent is

$$E(s_2|s_1, d = 0) = P(c = 0|d = 0)E(s_2|s_1, c = 0) + P(c = 1|d = 0)E(s_2|s_1, s_2 < -\frac{\beta_{01}}{\beta_{21}} - \frac{\beta_{11}}{\beta_{21}}s_1)$$

Define $\psi \equiv P(c = 1|d = 0)$, $\kappa_1 \equiv \frac{\sigma_1^2}{\sigma_2^2+\sigma_1^2\sigma_2^2+\sigma_1^2}$, $\kappa_2 \equiv \frac{\sigma_2^2}{\sigma_2^2+\sigma_1^2\sigma_2^2+\sigma_1^2}$, and $\omega = \sqrt{(1 + \sigma_2^2)(1 - \rho^2)}$. Then

$$E(s_2|s_1, d=0) = \frac{1-\psi}{1+\sigma_1^2}s_1 + \frac{\psi}{1+\sigma_1^2}s_1 - \psi\sqrt{(1+\sigma_2^2)(1-\rho^2)}\frac{\phi\left(\frac{-\frac{\beta_{01}}{\beta_{21}} - \frac{\beta_{11}}{\beta_{21}}s_1 - \frac{1}{1+\sigma_1^2}s_1}{\sqrt{(1+\sigma_2^2)(1-\rho^2)}}\right)}{\Phi\left(\frac{-\frac{\beta_{01}}{\beta_{21}} - \frac{\beta_{11}}{\beta_{21}}s_1 - \frac{1}{1+\sigma_1^2}s_1}{\sqrt{(1+\sigma_2^2)(1-\rho^2)}}\right)}$$

$$\cong -\psi(0.64\frac{\beta_{01}}{\beta_{21}} + 0.8\sqrt{(1+\sigma_2^2)(1-\rho^2)} + (\frac{1-0.64\psi}{1+\sigma_1^2} - 0.64\psi\frac{\beta_{11}}{\beta_{21}})s_1$$

When the job-seeker chooses not to send the letter the employer has to replace the observed value of s_2 with its conditional expectation $E(s_2|s_1, d=0)$

$$E(a|s_1, d=0) = (\beta_{00} + \beta_{01}) + (\beta_{10} + \beta_{11})s_1 + \beta_{21}E(s_2|s_1, d=0)$$

$$\cong -0.64\beta_{01}\psi - 0.8\beta_{21}\psi\sqrt{(1+\sigma_2^2)(1-\rho^2)} + \left(\frac{1-0.64\psi}{1+\sigma_1^2}\beta_{21} - 0.64\psi\beta_{11} + (\beta_{10} + \beta_{11})\right)s_1$$

So the coefficients of

$$E(a|s_1, d=0) = \beta_{00} + \beta_{10}s_1$$

$$\beta_{00} = -\frac{0.8\psi}{1-0.64\psi}\kappa_1$$

$$\beta_{10} = \frac{1}{1+\sigma_1^2}\kappa_1 + \kappa_2$$

Then

$$\beta_{01} = \frac{0.8\psi}{1 - 0.64\psi}\kappa_1$$

$$\beta_{11} = -\frac{1}{1 + \sigma_1^2}\kappa_1$$

$$\beta_{21} = \kappa_1$$

The perfect Bayesian equilibrium (PBE)[97] for this dynamic game of incomplete information is then that the job-seeker's decision to send the letter can be expressed as

$$d(s_1, s_2,) = c.1\left[s_2 - 1/(1 + \sigma_1^2)s_1 > -\frac{0.8\psi}{1 - 0.64\psi}\omega\right]$$

while the firm's hiring decision will be

$E(a|s_1, s_2, d) =$

$$1.\left[\frac{0.8\psi}{1 - 0.64\psi}\kappa_1\omega + -\frac{0.8\psi}{1 - 0.64\psi}\kappa_1\omega d + (\frac{1}{1 + \sigma_1^2}\kappa_1 + \kappa_2)s_1 - (\frac{1}{1 + \sigma_1^2}\kappa_1)s_1 d + \kappa_1 s_2 d > \theta\right]$$

[97]A PBE is a strategy profile and belief system that are sequentially rational and consistent. In our context, employers know the decision problem of the job seeker, who in turn knows that the hiring firm has this information. Neither firm nor job seeker can benefit by deviating from their strategy.

The share of applicants π who use of the reference letter in equilibrium ($P(c = 1)$) is

$$P(d(s_1, s_2,) = 1) = \pi.P(s_2 - \frac{1}{1 + \sigma_1^2} s_1 > -\frac{0.8\psi}{1 - 0.64\psi}\omega)$$

Simulation: How do effects change as letters become widely adopted?

Section 1.4 finds that increasing the number of letters from one to three does not change their effectiveness. However, this still presents a relatively small share of the applicant pool as employers report to typically receive about 50 applications per vacancy. While the experiment cannot create substantive variation in the share of applications submitted with reference letters, we can employ our conceptual framework (Section 1.2.2) to investigate how effects may change as reference letters become more widely adopted.

Simulation results presented in Figure A.8 illustrate the relationship between the share with access to letters (c) and the share using it in the job search (d). The share using the letter increases for two reasons: first, more high ability people with positive letters gain access and, second, the share using the letter conditional on having access to it $P(d = 1|c = 1)$ increases. One corollary is that the ability of the marginal job seeker using the letter decreases as more get access (results not shown). The intuition behind these results is that in an equilibrium with very few reference letters, the information employers can infer about the applicant's ability from not receiving a letter is limited. Workers who receive a negative letter can hide in the larger pool of job seekers without access to letters and will thus only accrue a small penalty, i.e. firms only slightly adjust their beliefs about ability downward.

Figure A.9 depicts how the firms' screening ability changes as more people get access to letters. Two effects emerge: the overall share of people hired $P(h = 1)$ increases suggesting that the letter has positive net employment effects and the average ability of the hired person $E(a|h = 1)$ increases.[98] Combined, these results suggest that the unraveling of the market (i.e. $P(c = 1) \rightarrow 1$) is desirable from a market efficiency perspective as it maximizes the information available to firms to identify the most able candidates.

[98] Pissarides, 1985 concludes that labor demand can change very rapidly as firms make hiring decision based on the *perceived* cost and benefits of future matches.

Figure A.8: Simulation: Reference Letter Access and Usage

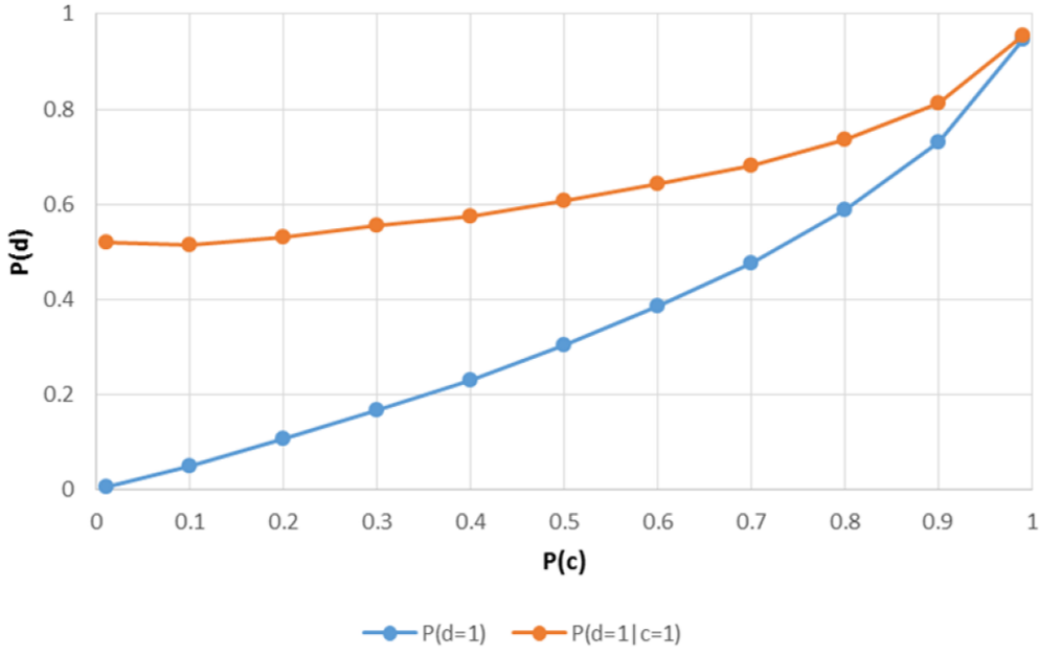

Note: The graph shows simulation results of the relationship between access to the letter (c) and using it (d)

Figure A.9: Simulation: Hiring and Ability

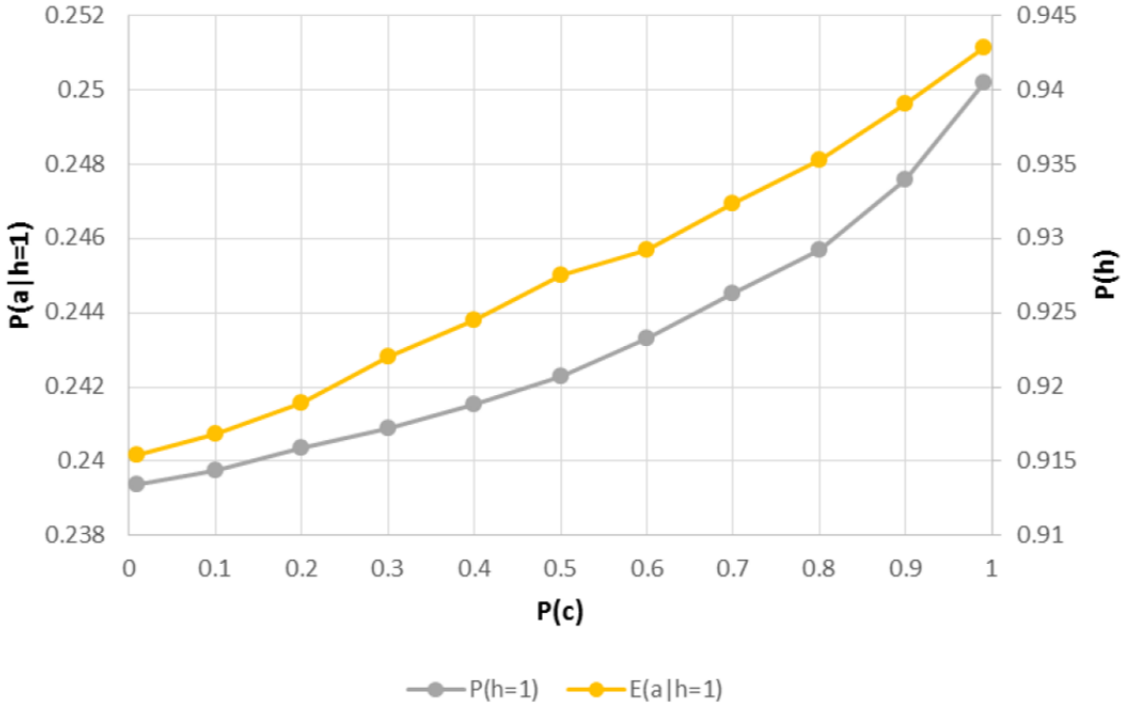

Note: The graph shows simulation results of the probability of hiring (h) and the expected ability (a) of hires.

B Appendix Chapter 2

Figures

Figure B.1: Action Plan Template

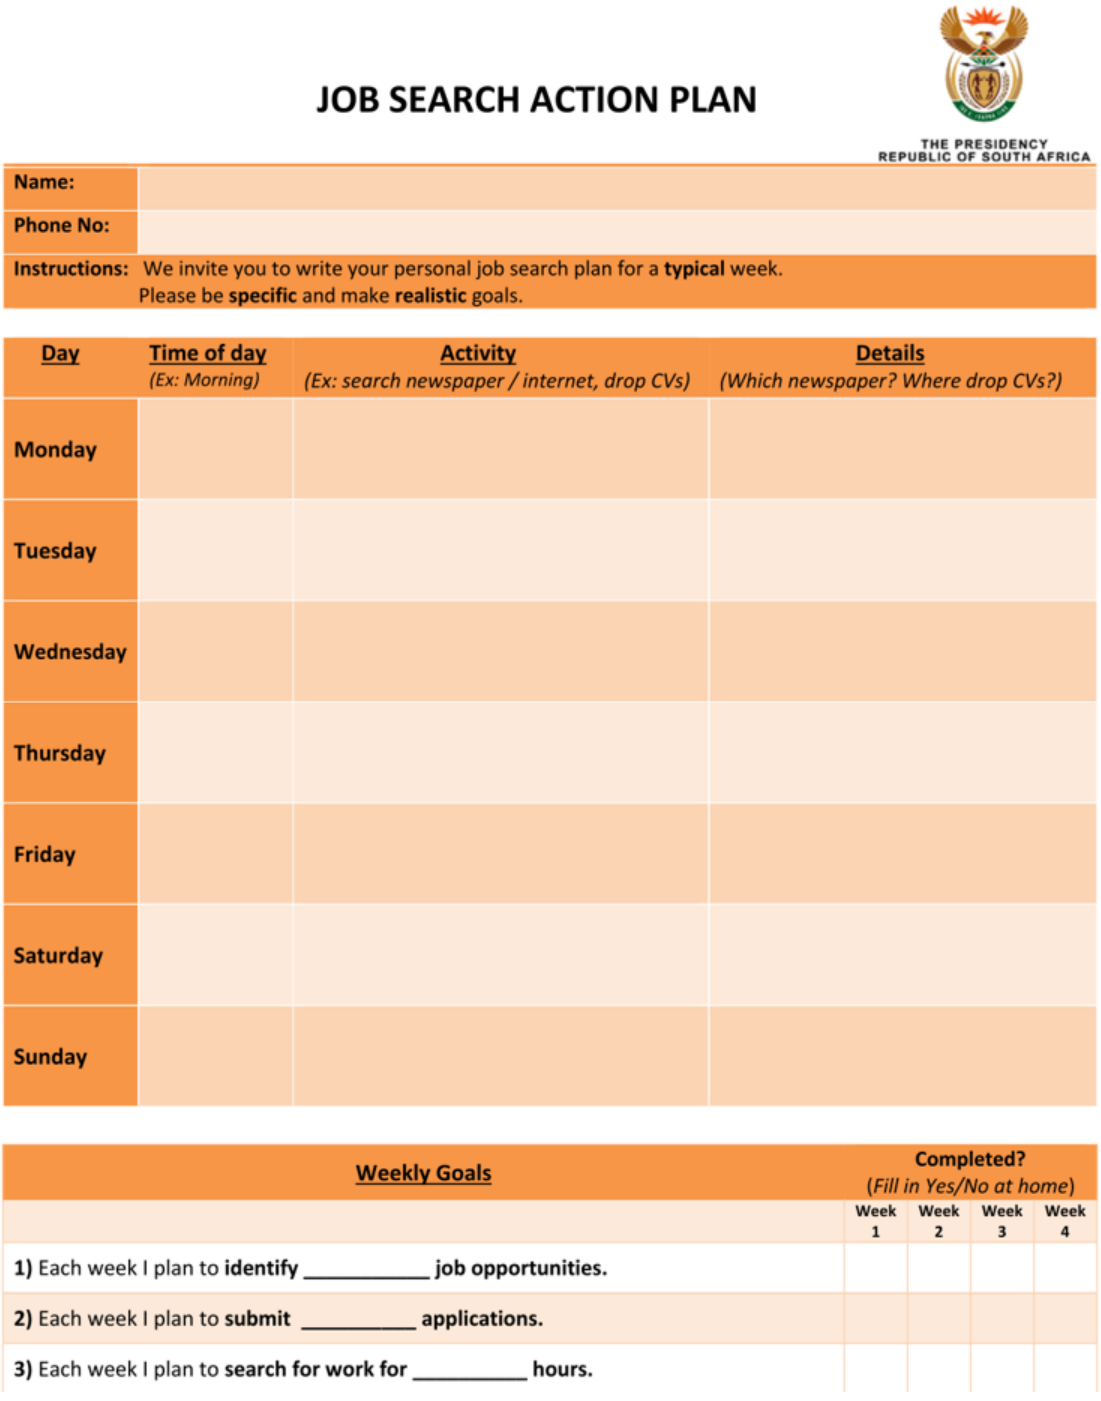

JOB SEARCH ACTION PLAN

THE PRESIDENCY
REPUBLIC OF SOUTH AFRICA

Name:	
Phone No:	
Instructions:	We invite you to write your personal job search plan for a **typical** week. Please be **specific** and make **realistic** goals.

Day	**Time of day** *(Ex: Morning)*	**Activity** *(Ex: search newspaper / internet, drop CVs)*	**Details** *(Which newspaper? Where drop CVs?)*
Monday			
Tuesday			
Wednesday			
Thursday			
Friday			
Saturday			
Sunday			

Weekly Goals	**Completed?** *(Fill in Yes/No at home)*			
	Week 1	**Week 2**	**Week 3**	**Week 4**
1) Each week I plan to **identify** _____ **job opportunities.**				
2) Each week I plan to **submit** _____ **applications.**				
3) Each week I plan to **search for work for** _____ **hours.**				

Tables

Table B.1: Attrition by Treatment

	Sample	Control	WSBasic	WSPlus	P-v. Basic vs. Control	P-v. Plus vs. Control
Attrition Wave 1	.05	.05	.04	.06	.52	.64
Attrition Wave 2	.15	.15	.15	.15	.91	.96

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Attrition rates are computed 5 weeks (Wave 1) and 12 weeks (Wave 2) after the intervention. The last two columns present p-value of tests of equal attrition.

Table B.2: Balance Test

	Control		Workshop		Action Plan		Peer		
	N	Mean	Mean	P-val.	Mean	P-val.	Mean	P-val.	F-Test
Age (years)	1097	26.38	26.71	.37	26.97	.12	27.1	.05	.1
Female	1097	.55	.5	.17	.48	.1	.49	.1	.1
Education (years)	1096	12.16	12.1	.58	12.08	.4	12.11	.63	.69
HH Size (adults)	1097	2.46	2.4	.72	2.29	.26	2.18	.08	.22
Employed HHMs	1097	1	1.01	.84	.98	.84	1	.98	.96
Searching HHMs	1096	.87	.8	.53	.77	.3	.7	.08	.28
1=married	1097	.06	.06	.86	.06	.99	.07	.44	.88
1=moved to Joburg	1097	.27	.3	.36	.34	.09	.32	.19	.16
Receiving UIF	1097	.07	.11	.12	.11	.13	.08	.69	.22
1=ever had job	1097	.77	.79	.63	.83	.07	.82	.19	.17
1=ever selfemployed	1097	.14	.14	.83	.14	.96	.19	.12	.48
Unemploym. (months)	870	13.93	12.96	.32	10.88	0	11.35	.01	0
Transport costs (BS)	1043	72.9	83.76	.15	64.68	.27	92.06	.02	.35
Nr Employed Friends	1097	1.8	2.1	.09	1.81	.96	1.89	.59	.19
Hrs/week, BL	1058	10.78	11.39	.47	12.34	.07	11.69	.28	.2
Apps/month, BL	1087	4.11	4.39	.5	4.42	.49	4.86	.11	.32
Responses/month, BL	1088	.44	.52	.29	.49	.5	.43	.84	.54
Interviews/month, BL	1087	.23	.22	.73	.3	.19	.23	.99	.55
Offers/month, BL	1088	0	0	.	0	.	0	.	.

Notes: Standard errors in parentheses. Tables report results of a regression that uses predited search days from an interval regression as explanatory variables.* p<0.10, ** p<0.05, *** p<0.01

Table B.3: Search Channel Usage (Interval Regression)

	Empl Agency	Drop CVs	Place Ad	Answer Ad	Search Online	Fam/Friends
	(1)	(2)	(3)	(4)	(5)	(6)
WS Basic	0.003	-0.176	0.244	-0.096	0.250	0.014
	(0.258)	(0.183)	(0.388)	(0.200)	(0.398)	(0.362)
WS Plus	0.770***	0.430***	0.320	0.570***	1.226***	-0.081
	(0.232)	(0.164)	(0.303)	(0.180)	(0.329)	(0.305)
Covariates	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1937	1936	1934	1927	1931	1926

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

C Appendix Chapter 3

Figures

Figure C.1: Job Website: Search Result

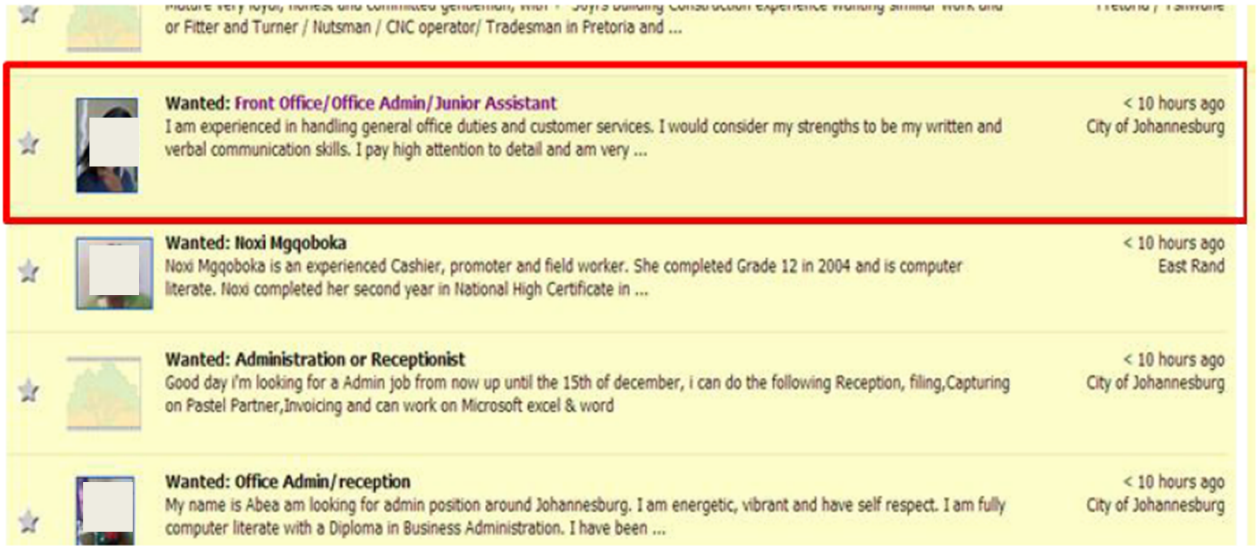

Figure C.2: Job Website: Profile Page

Figure C.3: Age Distribution: Job website vs. Labor Force Survey

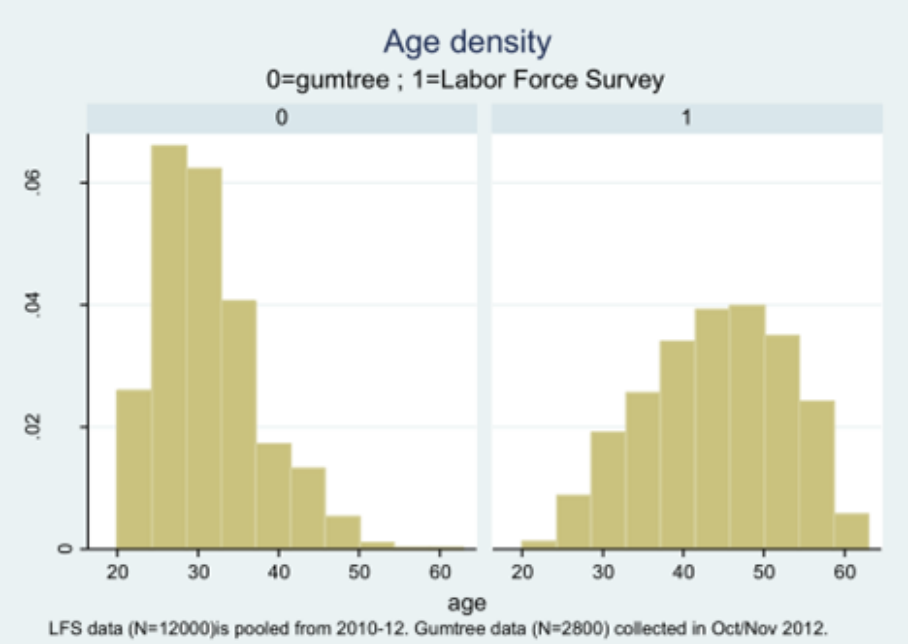

Tables

Table C.1: Job seeker characteristics by sector

			Sector			
	Mean	N	Housek	Nanny	General	p-value
Nr profile clicks	7.41	5341	6.61	9.66	7.43	0
South African	.035	5341	.03	.045	.04	.038
Malawian	.356	5341	.42	.15	.381	0
Zimbabwian	.189	5341	.21	.208	.101	0
No nation. info	.427	5341	.344	.599	.5	0
1=female	.429	5341	.515	.502	.065	0
1=male	.234	5341	.194	.128	.488	0
report age	.291	5341	.283	.353	.243	0
age (yr)	29.42	1552	30.20	28.33	28.31	0
report experience	.433	5341	.448	.399	.423	.014
experience (yr)	4.4	843	3.9	5.4	4.9	0
Reference	.296	5341	.335	.26	.207	0
Refer. phone nr	.014	5341	.023	0	0	0
Employer posts	.115	5341	.089	.165	.142	0
1=workpermit	.048	5341	.048	.049	.051	.914
Sleep in	.154	5341	.164	.214	.054	0
able to drive	.071	5341	.041	.062	.178	0
Certificate	.034	5341	.025	.073	.017	0
virtuout traits	.356	5341	.412	.284	.257	0
Nr words	31.1	5341	30.4	33.5	30.9	0
Picture	.087	5341	.045	.161	.138	0

Note: The table reports mean profile characteristics of job seeker profiles compared across sectors. P-values are reported of a test of equal means across sectors.

Table C.2: Robustness Test: Short vs Long profiles

	Full Sample		Housekeeper		Nanny		General Work	
	Long	Short	Long	Short	Long	Short	Long	Short
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Zimbabwe	-0.0924***	-0.0303	-0.0535	-0.0260	-0.139***	-0.0966	-0.0625	-0.0373
	(-3.62)	(-1.11)	(-1.57)	(-0.83)	(-2.82)	(-1.54)	(-0.71)	(-0.34)
Malawi	-0.144***	-0.0496**	-0.119***	-0.0522*	-0.195***	-0.0900	-0.0406	-0.0180
	(-6.39)	(-2.06)	(-4.17)	(-1.82)	(-3.70)	(-1.33)	(-0.60)	(-0.27)
R^2	0.303	0.161	0.254	0.156	0.328	0.245	0.259	0.176
N	2871	2384	1571	1609	774	340	526	435

t statistics in parentheses

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

Simple Model of Hiring

A simple model may help to clarify these three channels: assume that firms have the following utility function $U = U(Y, \lambda, \delta, c)$. Aside from the productivity of workers (Y), firms care about the fines (δ) they receive for hiring in undocumented immigrant with probability p which is a function of the legal status l. Last, employers may incur utility loss (λ) from interacting with foreigners. As evidence shows that South Africans' attitude towards foreigners varies across nationalities (Crush 2008), the distate is modeled to be a function of the applicants' country of origin $\lambda(d)$. In addition, the firm accrues cost c for reviewing each applicant. While I later consider a more realistic tournament model of hiring, for simplicity assume that there are n vacancies and n applicants. For each job applicant, the risk-neutral firm decides to hire iff:

$$E[Y|d, x] - \lambda(d) > c + \bar{w} + p(l)\delta \tag{16}$$

The condition simply states that the firm hires a native if the expected output is higher than the wage plus the screening costs c. For a foreign applicant, the expected profit must also exceed the expected fine and the distate from hiring a foreigner. For now, I assume that wages are fixed ($\bar{w}$) so firms strictly prefer to hire more productive workers. The model predicts that conditional on applicants' observable characteristics X, employers prefer to hire South Africans for three possible reasons:

1. Employers risk paying a fine if they are caught hiring an immigrant who is undocumented (i.e. $l = 0$).

2. Consistent with models of taste discrimination, employers may receive disutility from hiring and interacting with foreigners.

3. Consistent with models of statistical discrimination, employers may belief that South Africans are more productive with respect to unobservables e, i.e. $E[Y|x, d = SA] > E[Y|s, d = Foreign]$.

Employer Learning Model

To shed light on the nature of discrimination, I test a simple employer learning model based on Faber and Gibbons (Farber and Gibbons, 1996, FG) and Altonji and Pieret (Altonji and Pierret, 2001, AP). FG model productivity Y as a function of $\tilde{Y}(x, q, n, z)$. x captures information available to both the researcher and employer, q information available only to employers, z information only available to the researcher and n unobserved factors. Since I observe exactly the same information set as the employer at the time of the screening decision, the production function thus simplifies to $Y = \tilde{Y}(x_i, n_i)$.

Let's posit an additive separable production function in which a worker produces output $y = f(d)+g(x)+e$ with $e \sim N(0, \frac{1}{h_e})$ and let the ability y of each job seeker be a random draw from the population distribution of their nationality d. The employer observes the applicants' nationality d and forms beliefs $E[y|d \sim N(m_0(d), \frac{1}{h_0})]$. The employer then receives additional information x with $y|x \sim N(m_1, \frac{1}{h_1})$ which she uses to form the posterior belief $y|x, d \sim N(\frac{h_0 m_0 + h_1 m_1}{h_0 + h_1}, \frac{1}{h_0 + h_1})$.[99] It is straightforward to show that the posterior productivity belief is a weighted average of the information on nationality and other signals with the weights determined by the relative informativeness presented by the inverse of the variance of the error term.

This simple model of employer learning offers a test to distinguishing taste from statistical discrimination similar to AP (2001). Assume that the vector X consists of J potential

[99] I assume that h_i, which captures the inverse of the population variance σ_i^2, is constant, independent of d.

predictors of productivity Y such as age, experience or education, $X = (x_1, x_2, ..., x_J)$, that the job seeker may reveal on her profile.[100] Let's define $I = \frac{\sigma_x^2}{\sigma_e^2}$ as the relative informativeness of the observed set X relative to the unobserved error term e and assume that I monotonically increases in the number of signals j actually provided, i.e. each additional information increases the predictiveness of Y.

AP test whether the effect of the easy to observe variables (e.g. education, race) decrease as the employer learns additional information on the hard-to-observe variables over time. The equivalent test in this setting is to test if nationality becomes less important in the screening decision as job seekers provide more information. In a model of *statistical* discrimination,

$$\frac{\partial^2 P(Hire = 1|X, I, d)}{\partial d \partial I} < 0 \tag{17}$$

, i.e. the effect of nationality d on the probability of an applicant beeing hired decreases as additional information I is available (conditional on the actual content of the new information $X = x$). The intuition behind this prediction is straightforward: if firms use nationality as a proxy for productivity, then the importance of nationality should decrease as other predictors of productivity become available and are factored into employer beliefs. By contrast, if employers have the same productivity expectation of South Africans and foreigners but prefer hiring locals due to *taste*-based discrimination, we would expect the revelation of additional information to have no effect on the foreign-national gap in hiring decisions, i.e. $\frac{\partial^2 P(Hire=1|X,I,d)}{\partial d \partial I} = 0$.[101]

---

[100] Given that there is no standardized form to report information, job seekers may be strategic about what they reveal. For 201 job seekers in the domestic work sector, I observe *both* what is revealed online and detailed data from an anonymous survey. I find that there is no significant correlation between the reported age, years of experience, legal status or nationality with the decision to whether a person reveals this information in the classified ad. This could be explained by two factors: either job seekers do not know about the benefit of revealing positive characteristics and/or negative characteristics (e.g. immigrant status) is easily verifiable by employers so job seekers have an incentive to reveal it before occuring interview expenses.

[101] This test is distinct from related audit studies. Bertrand and Mullainathan (2004) and Oreopolous (2011) test how callback rates change for higher quality CVs. The quality of resumes is improved both by

Domestic Worker Phone Survey

Protocol: We collected phone numbers of all job seekers who posted in the domestic work category in Gauteng and the Western Cape between December 1st, 2015 and January 10th, 2016. Experienced phone surveyors called these people and explained that they are calling as part of a research project to "*understand the situation of domestic workers in South Africa.*" and that you "*will be asked about your job search and work history.*" Surveyors stressed that we are not offering employment and that the survey is completely volunary, fully anonymous, and would take 15 minutes to complete. "*As a thank you for participating in this research, you will receive 30 Rand Airtime, regardless of your responses and how many questions you choose answer.*" The compensation is about 20% of the daily income of domestic workers. It was paid via an airtime transfer to a phone number of their choice.

Selection: We attempted to call a total of 444 people of which we reached 343 (77.2%). Of these people we reached 303 (88.3%) agreed to participate in the survey. This is a remarkably high share compared to other phone surveys. To assess how selective this sample is, in particular with regard to nationatlity composition, I test whether the nationality information posted on the website is correlated with the probability of being reached or the decision to participate (conditional on being reached). Compared to people who do not post their nationality, foreigners are 0.3% more likely to be reached (p-value:0.96) and 2.7% more likely to participate (p-value: 0.54). South Africans are 3.1% more likely to be reached (p-value:0.29) and 3.3% more likely to participate (p-value:0.77). These results suggest that the sample is representative of the population of job seekers, at least with regard to nationality composition.

---

adding new information (e.g. additional certificates) *and* by changing the quality of signals provided (e.g. domestic vs. international work experience). The test in equation (17) by contrast looks at the effect of providing more signals (j) holding the quality of provided information (X) constant.