Investigating Hiring Frictions in Small Firms: Evidence from an Internet Platform-based Experiment

Pre-Analysis Plan

A. N. Fernando, N. Singh and G. Tourek

November 15, 2018

1 Introduction

1.1 Motivation

Rapidly modernizing cities in the developing world present new opportunities and challenges for small firms accustomed to traditional labor recruitment practices. For example, in India, e-commerce has created new opportunities for small firms providing the 'offline' delivery and logistics of the trade. Even as these firms look to adjust to rising demand, a pilot study we conducted shows a vast majority of them continue to rely on personal networks to recruit workers, driven in part by concerns regarding employee malfeasance. Such 'hiring frictions' lead to artificially small recruitment pools and may result in both lower worker quality and less stable matches. However, firms may be hesitant to hire outside of these networks as they provide credible information on workers that might be especially valuable in managing reputational and operational risks where employees directly handle merchandise and interact with clients.

We use a field experiment to assess whether an online job platform can improve the scope of recruitment pools available to these firms without compromising on the quality of information available on candidates. In a partnership with QuikrJobs, India's largest online portal for low and semi-skilled jobs, we implement interventions that expand candidate pools for firms on the online platform and provide verified information intended to mimic the screening function of personal networks in hiring. We then estimate the individual and joint impacts of these interventions on the hiring and retention of workers. In addition, our design will speak to whether effects are driven by worker 'types' or the 'actions' they take when hired.

Previous research on constraints to firm growth in developing countries has focused on the role of credit market imperfections (De Mel et al., 2008; Banerjee and Duflo, 2014) and management practices (Bloom et al., 2013). Our focus on demand-side labor market frictions as a constraint is relatively understudied in the literature; a few notable exceptions include work by De Mel et al. (2016) in Sri Lanka, Hardy and McCasland (2015) in Ghana, and Algan et al. (2018) in France. Our study will add to this literature on demand-side labor market frictions targeting a sample population that is of particular interest in developing country settings: small urban firms recruiting low and semi-skilled workers.

As mobile and internet penetration continue to grow across the developing world, it is inevitable that technology will shape traditional labor market institutions. Understanding how these small firms can leverage low-cost, scalable, technologies to address frictions specific to low and semi-skilled labor demand is of particular significance to urban poverty.

1.2 Research questions

Our main objective is to understand whether and how hiring frictions constrain firm growth and performance in developing economies, with a focus on the effects of expanding hiring pools and providing access to screening tools. Specifically, we seek to answer the following questions:

1. Does expanding the pool of potential workers available to firms lead to changes in hiring and retention of workers?

2. Do background verification services lead to changes in hiring and retention of workers?

3. How do these interventions interact to affect firm growth?
4. Are effects driven by changing in who gets hired ('selection') or by changing the incentives of hired workers ('moral hazard')?

2 Experiment

2.1 Treatments

The experiment offers two interventions to employers recruiting on the platform. The first intervention expands the recruitment pool available to an employer by providing access to 'premium' recruitment services on the QuikrJobs platform. This treatment increases the salience of a job advertisement via 'top-of-page' placement and 'gold'-tagged advertising (see figure 1) for a fixed time period. It also provides priority access to additional human resources support for employers.

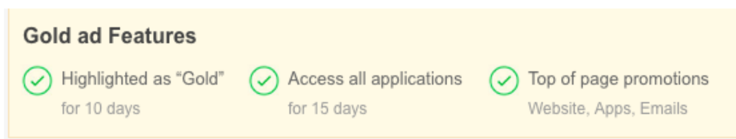

Figure 1: *Features provided under the 'premium' treatment*

The second intervention reveals verified information about applicant identity to employers. We first verify applicant identity using details they submit from government-issued ID during the application process and then randomly vary whether we reveal the outcome of their verification to the employer (see figure 2). An applicant's verification status is only privately revealed to the employer via a badge shown on the candidate's application within 24 hours of submission; that is, this information is not visible to other employers recruiting on the platform. Furthermore, we randomly offer some employers an additional type of background verification– employers receive an offer to verify the residential addresses of up to 3 applicants and must initiate the verification request. This verification is conducted through physical visits or postal checks by a third-party provider and takes an average of 5 days to complete.

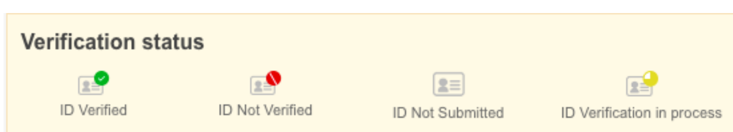

Figure 2: *List of verification outcomes visible to eligible employers*

2.2 Design

The experimental design cross-randomizes our two treatments— premium recruitment and background verification services. Our sample is split between those receiving the 'Regular (R)' service available at no-cost on the website and those receiving access to 'Premium (P)' recruitment services. Each employer in these groups is further assigned to one of three verification intensities for their pool of applicants: (i) 0%, verification status is never revealed; (ii) 50%, verification status is revealed for half of the applicants; and (iii) 100%, verification status is revealed for all applicants. In the 100% verification cells, we randomly offer the address verification service to a half of the employers. See figure 3 for a summary of this design.

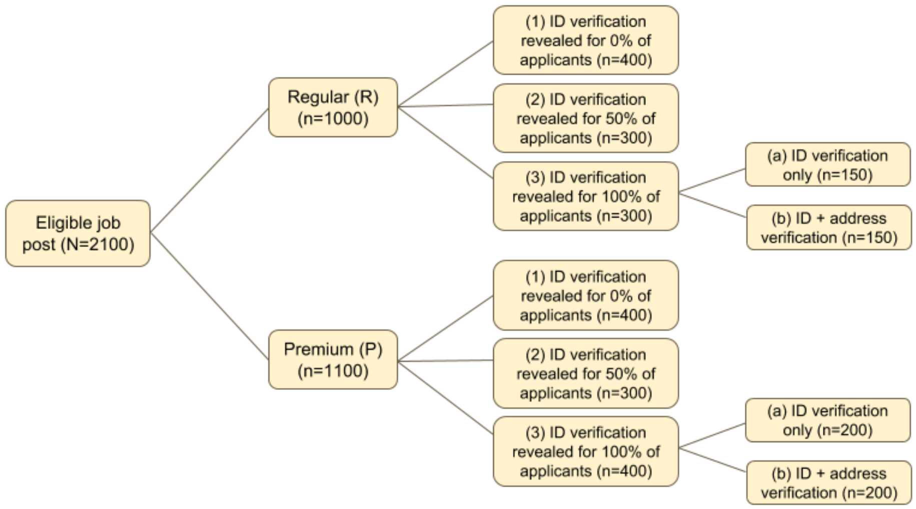

Figure 3: *Experimental design at the employer level with target sample sizes*

2.3 Sampling and Randomization

Our sampling frame is comprised of Bangalore-based firms posting vacancies in the sales, marketing and delivery job categories that have never purchased any recruitment services from the QuikrJobs platform. We restrict our randomization sample to only those firms that self-report having fewer than 50 employees. Screening takes place via online prompts during the job posting process and is verified in-person during the first round of surveying; large firms, employment agencies, and households are excluded for our sample. Power calculations support a target sample size of 2,100 employers to detect effects on hiring and retention from our interventions; we implement unequal group sizes, given in figure 3, based on our interest in specific comparisons (listed in section 4) and budget constraints.

Randomization occurs in two stages in real-time on the QuikrJobs platform. In the first stage, (eligible) employer job posts are assigned to one of our experimental groups (see figure 3). This assignment is stratified by job category, previous usage of the platform, and a firm size dummy for whether firm is below 10 employees. Then, based on this assignment, in the second stage, we randomize whether an applicant's verification status is revealed to the employer. The status of applicants is never revealed in the 0% cells and always revealed in the 100% cells. In the 50% verification cells, we randomize whether the status of the first applicant is revealed and then follow an alternating pattern.

3 Data and Outcomes

3.1 Data sources

We combine administrative data recorded by the QuikrJobs platform with survey data to understand the full scope of the experimental impacts. We survey employers four times—once at the start of the study and again at the 3, 6, and 12-month intervals. In addition, we administer cognitive (English, Math, and Raven's matrices) and non-cognitive tests (Big-5) via email/SMS to all jobseekers who apply to any employer in our sample. We also plan to survey a subsample of 3,000 jobseekers drawn randomly from our experimental groups 3 months after the interventions.

3.2 Outcomes

(i) Recruitment via platform (admin)

- Number of applications: Total number of jobseekers who applied to the vacancy
- Callback rate: Number of employer clicks to view applicant contact information/Number of applicants;
- Applicant quality: We will calculate average educational attainment and average years of experience as measures of applicant quality for each job post. For applicants who complete our cognitive tests, we will also be able to compute for each post average percentage of correct answers per subject and overall.
- Recruitment cost: Total amount spent on recruitment features available on the platform
- Number of new vacancies: Number of new job postings created by each employer on the platform

(ii) Recruitment from all methods (survey)

- Number of applications
- Callback rate: Number of applicants contacted by employer/Number of applicants
- Number of interviews/Interview rate
- Whether vacancy filled: An indicator for whether the firm filled the vacancy
- Whether hired applicant recruited through the platform: An indicator for whether a hired candidate was recruited through the QuikrJobs platform
- Number of new hires
- Time taken to fill vacancy: Number of days the vacancy was open
- Take-up of address verification: An indicator for whether an employer requested address verification of any applicants

(iii) Workforce (survey)

- Firm size: Number of workers employed at the firm, inclusive of all types of workers
- Fraction of employees who are [characteristic]: Number of workers who are [characteristic]/Number of workers. We will consider the following characteristics: migrant status; sex; education; religious affiliation; distance between residence and firm; contract type; caste
- Employee tenure (in years)
- Monthly income

(iv) Performance (survey)

- Revenues: Total revenues reported in the last complete month
- Costs: Total costs reported in the last complete month
- Losses due to employee misconduct: Value lost in monetary terms or working days due to employee misconduct

3.3 Managing attrition

We will employ the two techniques given below during data collection to minimize overall and differential attrition. Nevertheless, if we encounter issues, we will conduct robustness checks at the analysis stage using the methods described in section 4.2.

- Follow-up surveys: If a firm is missing in one round, our data collection plan offers multiple opportunities to track and obtain data.
- Randomized intensive tracking: We will follow a intensive tracking protocol for firms that do not complete surveys.

4 Empirical Analysis

4.1 Firm-level impacts

Our basic specification will estimate the intent-to-treat (ITT) effects of our interventions. We first list the set of these specifications, which map directly to our research questions. These regressions will include strata and survey round fixed effects, α_s and δ_t respectively, and cluster standard errors at the firm level. We will also include a pre-treatment measure of the dependent variable, Y_{i0}, as a control variable to soak up sampling variation.[1] Following this main analysis, we outline a plan for additional analysis to investigate mechanisms, exploit other variation generated by the experiment, or address any validity concerns.

[1] We will only be able to conduct our first survey with employers immediately after randomization occurs on the platform. While we expect no or limited hiring within the first few days of the vacancy being posted, we will exercise caution in using these variables as 'baseline' measures and will run our specifications both with and without this control.

4.1.1 Impacts of expanding the recruitment pool

Comparison: $R1, R2, R3$ *vs.* $P1, P2, P3$

We compute the pooled effects of premium recruitment services across all verification intensities by running the following specification.

$$Y_{it} = \alpha_s + \delta_t + \beta P_i + Y_{i0} + \epsilon_{it} \tag{1}$$

Here, P_i is an indicator variable for any job post that was randomly assigned to receive premium recruitment services. β gives us the ITT impact of premium recruitment service on the outcomes of interest.

4.1.2 Impacts of background verification services

Comparisons: (i) $R1, P1$ *vs.* $R2, P2, R3, P3$*; (ii)* $R1, P1$ *vs.* $R2, P2$ *vs.* $R3, P3$

The variation in this treatment will depend on whether applicants share information required for verification. Note that this treatment cannot be purchased separately on the platform.

Overall effects: We first compute the pooled effects of providing firms with (any) access to applicant verification.

$$Y_{it} = \alpha_s + \delta_t + \gamma V_i + Y_{i0} + \epsilon_{it} \tag{2}$$

Here, V_i is an indicator variable for any job post that was assigned to receive any applicant verification services (cells $R2$, $R3$, $P2$, $P3$). γ gives us the ITT impact of applicant verification services on the outcome of interest.

Effects by verification intensity: We separate the verification intensities in this next specification to understand whether impacts vary when employers receive different fractions of verified candidates.

$$Y_{it} = \alpha_s + \delta_t + \gamma_1(R2_i + P2_i) + \gamma_2(R3_i + P3_i) + Y_{i0} + \epsilon_{it} \tag{3}$$

$(R2_i + P2_i)$ is an indicator variable for posts receiving verification services for 50% of its applicants. Similarly, $R3_i + P3_i$ is an indicator variable for job posts assigned to receive verification services for 100% of its applicants. γ_1 and γ_2 give us the ITT impacts for the different verification intensities.

4.1.3 Joint impacts of premium and verification serices

Comparisons: (i) $R1$ *vs.* $P1$*; (ii)* $R1$ *vs.* $R2, R3$*; (iii)* $R1$ *vs.* $P1$ *vs.* $P2, P3$

To understand the joint impacts of the two interventions, we run the interacted regression given below.

$$Y_{it} = \alpha_s + \delta_t + \theta_1 P1_i + \theta_2(R2_i + R3_i) + \theta_3(P2_i + P3_i) + Y_{i0} + \epsilon_{it} \tag{4}$$

This regression allows us to understand the additional impact of verification or recruitment services on our outcomes of interest. We pool the verification intensities here as we are interested in overall impacts of access to any verified applicants. The omitted group in the regression is the set of employers in cell $R1$.

4.1.4 Importance of moral hazard concerns

Comparison: $R3a, P3a$ *vs.* $R3b, P3b$

We now compare employers who receive the ID verification status of all their applicants ($R3a, P3a$) to those employers who receive that and can additionally request address verifications for up to 3 applicants. Whether an employer takes up this offer gives us a measure of the importance of moral hazard concerns since having an applicant's verified address allows the employer to exert a stronger threat of punishment relative to only having an applicant's verified ID information. ψ_1 is our ITT estimate for the differential effect of the address verification offer on firm outcomes.

$$Y_{it} = \alpha_s + \delta_t + \psi_1(R3b_i + P3b_i) + Y_{i0} + \epsilon_{it} \tag{5}$$

Since employers choose the applicants for whom they would like to verify addresses, their choice may be correlated with unobserved applicant characteristics. Thus, we introduce selection in hiring and cannot simply conclude that any differential effects from the address verification treatment are driven by moral hazard concerns. We propose a strategy to address this selection in a future sub-treatment in section 4.4.

4.2 Further analysis

Time path of treatment effects: To detect whether treatment effects vary over time, we will interact our treatment indicators with time dummies, corresponding to our survey rounds. This would allow us to see, for example, whether any effects on firm size are short-lived or persistent. We may see short-lived effects if our treatments only change the timing of hires relative to the control group, which then catches up eventually.

Heterogeneity of treatment effects: We will explore heterogeneity in treatment effects by the stratifying variables (role, previous usage, and firm size dummy). This exploration will help us understand mechanisms that might be driving any effects we observe. For instance, our interventions may affect employers recruiting in delivery more than in marketing since delivery personnel have greater access to merchandise and interact directly with clients. Thus, trusting one's employees may be a more important factor in recruiting for delivery jobs and thus the marginal benefit of our background verification service would be greater for such employers. Apart from the stratifying variables, we also intend to explore heterogeneity by the median level of trust employers report in candidates they do not personally know (it was not feasible for us to obtain this information before randomization) and by whether the firm has previously hired outside of referral networks.

LATE impacts: We may have non-compliance to random assignment if $R1$ employers purchase platform-based products on their own. In this case, random assignment can be used as an instrument to estimate local average treatment effects. This analysis however will only be valid if our interventions do not affect the behavior of other employers on the platform. We discuss potential spillovers in section 4.3. Whether we use this strategy will thus depend on the degree of spillover we observe in order to ensure that we do not estimate spurious effects.

Candidate-level impacts: We will exploit experimental variation generated at the candidate-level to explore the effects of identity verification on jobseekers. One strategy considers cells where only 50% of applicant verifications are revealed. We compare outcomes for those candidates whose verifications are revealed with those whose aren't, relying only on variation within a job post since application decisions are not random. This strategy will allow us to understand spillover effects on candidates whose verifications are not revealed. We will also interact whether verification was revealed with the verification outcome (positive or negative, as seen in figure 2) to see how the type of information that is revealed affects outcomes.

A second strategy is to consider the full set of applications from candidates to employers in our sample. Since jobseekers may apply to numerous vacancies, we can consider how their probability of a callback changes when their verification status is revealed to an employer. We will use candidate fixed effects, holding constant selection effects from applying to particular jobs. This strategy is applicable for only those candidates who apply to multiple vacancies in our sample. We also plan to interact the indicator for whether an applicant's verification status with indicators for being in the 50% or the 100% verification cells. This gives us another measure of the spillover effect from background verification.

4.3 Threats to validity

Balance checks: We will compare our experimental groups pairwise, that is, we will test each treatment group with the $R1$ group as well as with each other using a linear regression model. These regressions will include strata-fixed effects. Along with pairwise comparisons, we will run (i) distributional tests to understand covariate overlap by experimental group distributions; and (ii) joint tests for whether our covariates predict treatment status. The data for these checks will come from two sources. The first set of variables will be details from job advertisements on the platform, including firm size, anticipated number of hires, and average profits. The second set of variables will be from our first survey of firms. This survey will take place soon after the vacancy is posted, but since this data will be collected after treatment assignment, we will exercise appropriate caution in the use of this data as 'baseline' measures.

Potential spillovers: Our sample is recruited based on the criteria listed in section 2.3. There are however other employers who recruit on the platform in our sample job categories and our interventions may lead to spillover effects on these other employers. Our first intervention may reduce the value of the platform for employers who purchase premium recruitment services on their own and for those employers who do not purchase any products but will now be listed further down the search page due to our intervention. Our second intervention may also affect who gets hired by employers outside our sample. The application process for out-of-sample employers, unlike all employers in our sample, does not require applicants to submit background verification information. As such, if certain types of applicants, for instance more dishonest applicants, prefer these out-of-sample employers over employers in our sample, our interventions may change the eventual hires of these out-of-sample employers.

These potential spillovers affect our ability to understand the broader welfare effects from the experiment. If these spillovers are large, then any positive effects from our interventions may just be due to poorer outcomes for out-of-sample platform firms, instead of net gains from the interventions.[2]

[2] We are less concerned with the displacement effects on 'off-platform' firms because the platform represents a

Our experimental design however provides a strategy to account for and address spillovers. This strategy relies on exploiting natural variation generated from our sample being realized over a period of a few months. We can use this variation to generate a measure of treatment exposure experienced by out-of-sample firms and investigate how variation in this exposure changes recruitment outcomes using administrative platform data. The same variation can also help us understand how the effects of differential exposure to sample firms on the platform changes jobseeker behavior. Finally, in practice, our sampling strategy may also mean that spillovers are limited because, at any given point in time, our sample job posts will likely account for a small fraction of the total job posts in a given category. Using platform data, we will be able to compute this fraction at high frequency over the duration of our interventions on the platform.

Attrition: We may observe attrition in spite of our field protocols. If so, we will conduct a number of additional tests. First, we will check overall rates of attrition and assess whether our attrition is differential across experimental groups. While high rates of attrition affects our power in detecting treatment effects, it does not affect the internal validity of our experiment as long as the attrition is not differential across groups. For the set of attriters, we will also examine whether pre-treatment characteristics are the same on average across groups. Finally, we will also check the robustness of our results to using predicted attrition weights, which we calculate by predicting the probability of refusal given treatment status and covariates.

4.4 Additional treatments under consideration

Based on early results from platform and survey data, we will consider additional treatments in order to more carefully understand mechanisms or spillovers. We list the strategies and the reasoning for these additional treatments here, but will specify the exact design and analysis plan for evaluating these treatments before we conduct the randomization.

First, if we observe large effects on hiring of platform candidates from the verification treatment, we will introduce a sub-treatment to understand the role of moral hazard directly. The take-up of the address verification offer by employers in cells $R3b$ and $P3b$ only provides us with a measure of the severity of employer concern regarding employee moral hazard. It cannot help us to disentangle whether the impacts of ID verification are the result of a change in the type of worker hired or a reduction in employee moral hazard; this is because we introduce selection into the hiring process by allowing the employer to choose the applicants for whom address verification is conducted. To distinguish cleanly between the two explanations, we will thus randomly offer employee address verification at no-cost to a subset of employers only *after* they have hired candidates from the platform. Changing the timing of this offer allows us to focus exclusively on the moral hazard component since employers will have already made their hiring decisions.

Second, as discussed in section 4.3, our design may lead to spillovers. Depending on the severity of this spillover, in a future set of treatments, we will adopt a more direct approach to estimating these spillovers by randomizing the fraction of employers who are eligible to be part of the study over time and across job categories.

small fraction of the local labor market. As such, any effects would likely be negligble.

5 Timeline

The anticipated timeline of project activities is given below.

1. Interventions + First firm survey: December 2018 - February 2019
2. Applicant Tests: December 2018 - February 2019
3. Second firm survey: April - May 2019
4. Third firm survey: July - August 2019
5. Working paper draft: October 2019
6. Fourth firm survey: December 2019 - January 2020
7. Final draft: May 2020

References

Algan, Y., Crepon, B., and Glover, D. (2018). The value of a vacancy: Evidence from a randomized evaluation with local employment agencies.

Banerjee, A. V. and Duflo, E. (2014). Do firms want to borrow more? testing credit constraints using a directed lending program. *Review of Economic Studies*, 81(2):572–607.

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013). Does management matter? evidence from india. *The Quarterly Journal of Economics*, 128(1):1–51.

De Mel, S., McKenzie, D., and Woodruff, C. (2008). Returns to capital in microenterprises: evidence from a field experiment. *The quarterly journal of Economics*, 123(4):1329–1372.

De Mel, S., McKenzie, D., and Woodruff, C. (2016). Labor drops: experimental evidence on the return to additional labor in microenterprises.

Hardy, M. and McCasland, J. (2015). Are small firms labor constrained? experimental evidence from ghana. *Unpublished*.