Expected Discrimination and Job Search*

Deivis Angeli [JMP]† Ieda Matavelli‡ Fernando Secco†

April 25th, 2024. See the latest version here.

Abstract

The impacts of labor market discrimination depend not only on whether employers discriminate, but also on jobseekers' responses to expected discrimination. To study these responses, we ran a set of field experiments with over 2,000 jobseekers in Rio de Janeiro's favelas, where most jobseekers overestimate anti-favela discrimination, as we measure it in an ancillary study. Jobseekers who were randomly told that their interviewer would know their name *and address* believed that their interview performance was 0.17SD worse than those who were told that the interviewer would only know their name. Focusing on jobseekers who expected at-or-above median discrimination, we find that not only did they believe that they performed worse when told that interviewers would know their addresses, their interviewers also rated them worse by 0.22SD. Removing the need to declare an address at the application stage increases interview attendance *only* for white jobseekers, likely because they can pass for non-favela residents and ignore racial discrimination. Our findings show that expected discrimination can create inefficiencies in matching, especially through effects on interview performance, and that correlated sources of discrimination like race mediate these effects.

*We are grateful for invaluable guidance from Matt Lowe, Siwan Anderson, Jamie McCasland, and Munir Squires. Beatriz Morgado Marcoje has provided unrivaled research assistance. We also thank Mackenzie Alston, Nava Ashraf, Leonardo Bursztyn, Claudio Ferraz, Ro'ee Levy, Rogério Santarrosa, Heather Sarsons, Devin Pope, Nathan Nunn, Chris Roth, Colin Sullivan, and seminar participants at the Vancouver School of Economics for their comments. This research was undertaken thanks to funding provided by the Canada Excellence Research Chairs program awarded to Dr Erik Snowberg in Data-Intensive Methods in Economics, by the Center for Effective Global Action (CEGA) through its Psychology and Economics of Poverty (PEP) Initiative, and by the J-PAL LAC Jobs and Opportunity Initiative Brazil. This study was approved by UBC's Behavioural Research Ethics Board (H22-03418), Insper's Research Ethics Committee (Opinion N. 281/2023), and pre-registered in the AEA RCT Registry (AEARCTR-0009359).

†University of British Columbia; deivisangeli@gmail.com, fernandoseccoluce@gmail.com

‡University of New South Wales; i.matavelli@unsw.edu.au

1 Introduction

Employers discriminate along many dimensions, including race, sexuality, and criminal history (Neumark, 2018; Rich, 2014; Riach and Rich, 2002). While quantifying employer discrimination reveals its pervasiveness, it may not tell us much about its ultimate effects, which depend on equilibrium factors. In particular, the equilibrium effects also depend on jobseekers' reactions to expected discrimination (i.e., their expectations about receiving different treatment because of some "irrelevant" characteristic). Jobseekers may hold miscalibrated beliefs, and beliefs about discrimination can even become self-fulfilling (Coate and Loury, 1993). Miscalibration may be problematic if individuals overestimate discrimination, becoming discouraged or too nervous to give their best performance at an interview. Beliefs may be self-perpetuating if they make some group appear worse on average to recruiters (e.g., discouraged applicants may invest less in their applications), allowing initial misperceptions to evolve into actuality. In this fashion, expected discrimination can also matter for policy. For example, if jobseekers overestimate discrimination, it might be desirable to disseminate information on actual discrimination rates, and for employers to credibly signal commitment to anti-discrimination policies. To investigate these ideas, we study the the effects of expected discrimination on job-seeking behaviors.

We study the effects of expected anti-favela discrimination on jobseekers in Rio de Janeiro, Brazil, where about 1.5 million people, 22% of the city's population, live in favelas. These neighborhoods are urban slums, where criminal organizations typically hold a monopoly over violence. Their residents are more likely to be migrants, less educated, and poorer than non-favela residents. There is also a strong correlation between race and favela residence, with the white population being a majority outside but only 33% inside favelas. The derogatory term "favelado", meaning "slummed", is widely used. In this context, most firms collect home address information from applicants, and, while this is meant to let firms to estimate commute length, it can also be used for discrimination regardless of distance to work. Jobseekers understand that discrimination as stemming from various factors: over 60% of jobseekers surveyed in this study mention violent police raids, racial and cultural prejudice, antipathy for favela residents, and fear of crime and violence as important reasons why firms avoid hiring people from favelas.

While our main goal is to understand how expected discrimination affects job application and interview performance, studying *anti-favela* discrimination also sheds light on two other topics. First, address-based discrimination may be relevant in perpetuating poverty traps in many contexts. Almost one billion people live in urban slums (UN, 2016), and even in developed countries, we see urban divides (e.g., public housing projects in the US). Can expected discrimination play a role in perpetuating such divides? Second, we argue that policymakers should not only consider the fact that people expect discrimination along some dimensions, but also how those dimensions interact. We document that race, a correlated source of discrimination, is an important determinant of the effects of expected anti-favela discrimination, and that policies that ignore race can lead to unintended consequences.

We ran three interconnected field experiments with over 2,000 favela jobseekers. We partner with a large cosmetics company to advertise and observe favela jobseekers applying and interviewing for real sales jobs. Jobseekers enter our study's pipeline through our door-to-door baseline survey in favelas, which shows that about 87% of jobseekers overestimate anti-favela discrimination – as measured in an ancillary audit study we ran by submitting 1,400 job applications with random home addresses. Our main strategy to randomize anti-favela discrimination emulates a address blinding policy, reducing expected address visibility. We use this strategy in two experiments, one varying whether one needs to declare their (favela) address in the job application form, and one varying whether the candidate may expect their interviewer to know their address. We keep everything else fixed on the side of the hiring firm (e.g., interviewers did not learn the candidates' addresses before evaluating them), so any effects should come from supply-side beliefs. Leveraging the fact that most jobseekers overestimate anti-favela discrimination in callbacks, we use a third experiment to estimate the effects of an information-provision policy (revealing our audit study findings) on job application decisions. Hence, by studying both application decisions and interview performance, and using two different strategies to shift expected anti-favela discrimination, our study provides a comprehensive picture of how expected discrimination affects jobseekers' behaviors. See Figure 1 for a simplified scheme showing how the experiments fit together.

Jobseekers enter our study's pipeline by answering a door-to-door survey, in which we gather baseline data. The key baseline variable is our incentivized measure of expected discrimination. We explain that we were also running an audit study, i.e., that we created fictitious worker profiles, randomized their addresses, and used those profiles to submit applications to 700 sales jobs in Rio, sending two different-profile applications to each. Then, we ask respondents to guess our findings, paying them based on accuracy. Over 85% predict anti-favela discrimination, while about 60% predict that having a favela address would cause callback rates to drop 50% or more. In the audit, we find no statistically significant discrimination – the callback rates were 19.3 and 19.6% for favela and non-favela résumés (p=0.38 for the difference). The absence of a statistically significant differences in callback rates does not imply there is no anti-favela discrimination at all, as audit studies are imperfect tools (Kessler et al., 2019; Neumark, 2018), nevertheless, the gap between prediction and reality shows that favela jobseekers are too pessimistic about address-based discrimination in callbacks.

To measure how jobseekers respond to expected discrimination in the field, we set up an HR firm that advertised real sales job opportunities in a large cosmetics firm.[1] At the beginning of the door-to-door survey, after some background questions, jobseekers could agree to share their professional details with this HR firm (described as a partner in the study). Within the next few days, the HR firm texted the jobseeker with an invitation to apply. Then, the HR firm invited applicants for interviews at its office in Downtown Rio. We used this structure to run three field experiments: two on the effects of expected discrimination on the application decisions and one focusing on interview performance.

In the Address Omission Experiment (N=1,303), we manipulate expected address visibility by randomizing the content of the application invite message and the online application form. In our main treatment condition, *Address Omission*, the text message stated that address information was unnecessary at that stage, and the form did not mention address at all. In our *Status Quo* condition, the text message listed the home address as necessary information for applying, and people need to fill it in. The address requirement does not affect average application rates: 42.7% of jobseekers applied in *Status Quo*, and 41% in *Address Omission* (p=0.62 for the difference). Considering all the invited applicants, 19.3% of those in *Status Quo* show up for the interview, and 19.8% in the *Address Omission* (p=0.64 for the difference).

One explanation for the null effect is that people can easily "pass" as non-favela residents, e.g., by declaring a different neighborhood or a relative's address in the *Status Quo* condition. That address obfuscation behavior was one of the anecdotes motivating this study, and, indeed, 28% of the *Status Quo* applicants obfuscated address. To understand the importance of obfuscation, we included an additional experimental condition that prevented it. This third condition, *Known Address*, was the same as *Status Quo* except that the online application form already contained the applicant's home address, and applicants just needed to double-check it. Nevertheless, *Known Address* generated application rates similar to the other two conditions. Even conditioning on the subsample who predicted at-or-above median discrimination in the audit study, there are no effects of changing expected address visibility on application decisions.

It is also unlikely that a failure to shift expected anti-favela discrimination is behind the null effects, for two reasons. First, reducing expected address visibility increases interview show up in the white subsample (one of our four pre-registered heterogeneity cuts), which is consistent with white jobseekers expecting to be able to pass as non-favela residents while non-white jobseekers expect discrimination either way – because their race will be eventually visible and

[1] These are not strongly gendered jobs: in our study, the application rate for males was 37%, and for non-males, it was 44%.

because observers will associate their race with residing in favelas. Second, in the Information Experiment, we verify that shifting beliefs about expected discrimination is not sufficient for having average effects on application rates – which will be our overall conclusion.

In the Information Experiment, we manipulated expected *market-level* discrimination using three experimental conditions: i) *No Info*, ii) *Favela Info* (revealing the audit study callback rate for a favela), and iii) *Full Info* (revealing that favela and non-favela callback rates were the same). *Full Info* reveals both the discrimination rate and callback level, so *Favela Info* works as an alternative control condition, holding constant the knowledge of the favela callback level. Providing information has a first-stage effect on expected anti-favela discrimination, as it affects (incentivized) posterior beliefs about how often the partner HR firm would callback applicants from inside or outside a favela. Regardless, jobseekers in the three conditions make it to the interview stage at the same rate of about 20%. Hence, the Information Experiment also suggests that average application rates are inelastic to expected anti-favela discrimination. Nevertheless, the racial heterogeneity now appears flipped, with white jobseekers being relatively discouraged after learning that we found no discrimination in callbacks for sales jobs. The most natural interpretation we can offer for this reversal is that low anti-favela discrimination implies low racial discrimination and low returns for passing. These would then imply more opportunities (mainly for non-white jobseekers), but also more competition for white jobseekers.

We ran the Interview Experiment (N=422, out of the $\approx$2,200 invited to apply) in an office staffed with one receptionist and up to two interviewers. On arrival, the receptionist asked jobseekers to confirm their name, date of birth, and address, then told them to wait. Then, the receptionist told the jobseeker that the interviewer was ready, and that, to keep the process objective, "the interviewer will only know your name" (*Name-Only* condition) or "your name and address" (*Name-and-Address*). The two conditions differ only by two words: "and address". Interviewers followed a script, did not know about the randomization, and learned about the jobseekers' neighborhoods of origin only after the end of the experiment. So, any treatment effects must be triggered by changes in the interviewees' behaviors or beliefs.[2]

Our main interview performance measures aggregate the interviewers' and interviewees' evaluations. After each interview, interviewers coded on 0–10 scales, i) how well the interviewee performed overall, ii) how nervous the interviewee was, and iii) how professionally the interviewee behaved. Candidates filled out a feedback form at the reception desk before leaving, including self-assessments for the same three dimensions. To maximize statistical power and reduce the risk of multiple hypothesis testing, we construct an inverse-covariance-weighted

[2]Note also that our manipulation rules out self-signaling effects (e.g., that reminding candidate of their address leads them to update negatively about their skill), since all jobseekers are asked about their address before treatment.

index of impressions for the interviewers and for the interviewee (Anderson, 2008). As our primary aggregate measure, we average the two.

Hearing that the interviewer will only know one's name increases the aggregate performance index by 0.13SD (p=0.03). Put another way, expecting address to be known (the status quo) decreases performance by 0.13SD. Unbundling the aggregate index, we see stronger effects on the self-assessment index (0.17SD, p<0.01). The effect size on the interviewer's evaluation index is 0.09SD (in the same direction), but it is not statistically significant nor different from the effects on self-assessment. Focusing on those who expected at-or-above-median discrimination at baseline, we see that increasing expected address visibility has a statistically significant negative effect of 0.22SD on the interviewer's evaluation index, consistent with high expected discrimination actually damaging interviewer-assessed performance. Hence, we find evidence of a self-fulfilling prophecy, at least in the narrow sense that if a jobseeker expects a worse evaluation due to their address, they indeed get one, even when interviewers can not discriminate.

While the Interview Experiment was not designed to pin down exactly which mechanism leads to impaired performance, our data is consistent with the idea that a mix of stress and heightened stakes leads jobseekers to "choke under pressure". First, it fits with the broad pattern of null average effects of expected discrimination on application decisions (a lower pressure private decision) and negative effects on interview performance (a setting with higher stakes and face-to-face interactions). Second, those who expected high discrimination reported to feel much more nervous when their addresses were known, consistent with an account of perceived discrimination or injustice leading to stress (e.g., Berger and Sarnyai (2015)). Third, we see evidence that jobseekers find it harder to be strategic at the interview office: i) they obfuscate their addresses less often at the office than on the the application form, and ii) expected address visibility has a significant negative average effect on self-reported professional behavior, suggesting that jobseekers do not control their behavior as well when address is visible. While other mechanisms such as reducing perceived returns to effort or overcompensating (e.g., by exaggerating qualities or acting too formally) are possible, it is not obvious that these would lead to the same empirical patterns. For instance, lower perceived returns to effort should not cause stress responses, and overcompensating should not cause lower professionalism in *Name-and-Address*.

Similar to when we reduced expected address visibility in the application procedure, we also find that reducing expected address visibility at the interview is better for white jobseekers. Being told that the interviewer only knows one's name increases performance by 0.31SD for white jobseekers, while only increasing performance by a non-significant 0.07SD for non-white candidates (p=0.09 when testing effect difference). Again, this is consistent with white jobseekers being more able to pass as non-favela residents, and not having to deal with race discrimination.[3] Hence, we find robust evidence that making a firm blind to address may exacerbate the inequalities stemming from the different but correlated race stigma.

Overall, we provide evidence that expected discrimination can exacerbate the impacts of employer discrimination on jobseekers, mainly by hurting interview performance, which gives rise to a type of self-fulfilling prophecy. Also, our findings that different policies tackling expected anti-favela discrimination generate race-heterogeneous effects suggest that policymakers should be mindful of overlapping perceived disadvantages, while keeping in mind that correlated sources of discrimination can interact in counterintuitive ways.

While many experiments measure whether agents discriminate in the labor market (see Neumark 2018; Rich 2014; Riach and Rich 2002 for reviews), the supply side has received much less experimental attention. In response, our field experiments are the first focusing on estimating the effects of expected discrimination.[4]

Closest to our study are three field experiments that experimentally varied the language in job ads (Del Carpio and Fujiwara, 2023; Burn et al., 2023) or how the selection procedure is described (Avery et al., 2023). These studies find effects on the composition of the applicant pool which could be explained by expected discrimination, but can not rule out other channels. For instance, the non-gendered (as opposed to gendered) job ads in Del Carpio and Fujiwara (2023) also signal work-life balance and an inclusive culture, which can appeal differently to females. Two main design elements allow us identify the effects of expected discrimination more sharply. First, we elicit beliefs about discrimination at baseline, allowing us to estimate whether expected discrimination predicts effect sizes. Second, we designed our experiments to vary only i) expected address visibility, keeping job desirability and other factors as constant as possible, or ii) expected market-level discrimination, which is not subject to job-level confounders.

We build on two lab studies that test whether jobseekers change how they present themselves in response to expected discrimination. Kang et al. (2016) shows that non-white college

[3] A Bayesian should assign a random white person in Rio a 13% chance of residing in a favela, and twice that chance for a non-white person. Hence, if candidates are careful not to hint at their home address by revealing information directly or through how they speak, a Bayesian interviewer should not guess that they are favela residents. Nevertheless, even a moderate probability of being from a favela may lead to discrimination, and a literature on stereotypes (see Bordalo et al. (2016)) suggests that observers exaggerate group differences, making this calculus even more favorable for white jobseekers. In interviews, only 4% of all jobseekers directly revealed that they were favela residents.

[4] A few observational studies find evidence consistent with expected discrimination affecting jobseekers. For instance, people who expect to be stereotyped might try to give clearer signals of their productivity (Lepage et al., 2022; Dickerson et al., 2022; Lang and Manove, 2011), or hide a stigma even when it is costly (Agüero et al., 2023). Closest to this study, Pager and Pedulla (2015) uses administrative and survey data, to show that Black jobseekers cast wider nets in their job searches and that breadth correlates with past experiences with discrimination. See also Kuhn and Shen (2023).

students craft "whitened" résumés (e.g., listing a Western name or omitting job experiences that could reveal ethnicity) but decrease the use of such strategies when crafting a résumé for a pro-diversity employer. Charness et al. 2020 finds that female college students are less likely to pick gender-matching avatars in a virtual labor market when competing for a male-dominated task. We go beyond by studying job applications and obfuscation strategies in the field.[5]

While it is possible that employers discriminate more at the interview than in callbacks (Quillian et al., 2020), interview performance has received little attention in the study of discrimination (and in labor economics as a whole). The main finding in Goldin and Rouse (2000) hints at the potential relevance of expected discrimination for interview performance: female hiring increases after orchestras adopt "blind" auditions. That effect could be both because evaluators lose the ability to discriminate and because females might perform better music knowing that they will be evaluated only on merit. While rarely discussed, our findings suggest that there is some merit to the latter account.[6]

We contribute to the study of discrimination as a self-fulfilling prophecy. In theory, even with no differences in initial endowments, pessimistic beliefs about returns to investment can make a group of workers acquire less human capital in response to expected discrimination, justifying statistical discrimination (Coate and Loury, 1993; Lundberg and Startz, 1983). Glover et al. (2017) shows that a similar kind of self-fulfilling prophecy can stem from managers' beliefs, which leads them to exert less effort in supervising minority cashiers over a trial period, making them less productive and less likely to be hired. We show how anticipated discrimination can also generate a kind of self-fulfilling prophecy in the matching process, exclusively through jobseeker's beliefs (as we hold the HR firm's actions constant).

Finally, our experiments randomizing expected address visibility connect with a literature on stereotype threat, which is the idea that when people feel at risk of confirming some negative stereotype, they may perform worse and confirm that stereotype (Steele and Aronson, 1995). That literature overwhelmingly considers test performance or other laboratory outcomes (Spencer et al., 2016; Liu et al., 2021), and we provide evidence that it can be relevant in a high-stakes job market context.

[5] There is also lab-in-the-field evidence that expected discrimination may affect on-the-job outcomes like retention (Ruebeck, 2024) and productivity (Hoff and Pandey, 2006). See also Fryer et al. (2005) for a classroom game using the Coate and Loury (1993) framework, and Aksoy et al. (2023) for an experiment on anticipated discrimination against LGBTQ+ supporters in the context of prosocial behavior.

[6] In a study with students, Word et al. (1974) provides a thought-provoking study of how even non-verbal interviewer cues triggered by a racial mismatch between interviewer and interviewees can lead to worse interview performance. The effects of expected discrimination on the job can also be important.

2 Context, Sample, and Misperceived Discrimination

2.1 Favelas in Rio de Janeiro

Brazilian favelas are areas of dense informal settlements. In Rio de Janeiro, the state has been unable to hold the monopoly of violence over favelas, which are home to 1.5 million people (one-fifth of the population). According to the 2010 Census, 66% of favela households had a per capita income of one minimum wage (≈10 USD/day) or less. Outside the favela, that rate is 30%, and per capita income is 3.5 times larger. Favela residents are also less likely to be literate (84% are literate inside favelas, 92% outside them), to have completed high school or an advanced degree, or to self-identify as white (33% in favelas and 57% outside).

Jobseekers in our study lived in Maré, Manguinhos, or Jacarezinho, three large adjacent favelas in Rio, home to about 200,000 people. These neighborhoods grew to their current boundaries without proper urban planning or public services, and now are part of a contiguous metropolitan area, sharing borders with other favelas and regular "asphalt" neighborhoods. We conducted most of our fieldwork in Maré, which is the most populous favela in Rio.

There are limited formal work opportunities in favelas. For instance, according to a census of Maré's Businesses conducted by a local NGO from 2011 to 2013, 75% of these businesses were entirely informal. In total, they employed only 9% of the favela's working-age population (REDES, 2014). Hence, most favela jobseekers must look outside for jobs. Formal jobs are specially attractive, since they provide more benefits and stability.

Residents in all three favelas are regularly exposed to violence or its imminent risk. In Maré, three criminal groups – two of which exploit the illegal drug market, and another working mainly as an extortion racket – hold the monopoly of violence. Criminal groups were also present in the two other favelas during our fieldwork, but police was also somewhat present.[7] Over our five months of fieldwork, police raids interrupted our survey activities 14 times. These police raids are generally unpredictable and violent. During a raid, favela residents will take refuge at their homes to avoid the crossfire. Workers miss work days, businesses shut their doors, and communications are hampered (internet or telephone). It is typically unclear when a police raid ends, so disruptions may persist for several days.

When there is no police raid in progress, favela residents can typically go in and out without issues. Some may work in the asphalt neighborhoods adjacent to their favela or commute to wealthier areas of the city. Commuting to these richer areas (e.g., Rio's Downtown or South

[7] See Lessing (2021) for a conceptualization of the symbiotic interaction of such criminal groups and the state. See also Monteiro et al. (2022) for an empirical account discussing the economic trade-off these gangs face, and Barnes (2022) for an ethnographic account of how gangs have responded to state action in recent years.

Zone) using public transportation may take 30 to 90 minutes. The Downtown office of our HR firm, where we held interviews, was within a 50-minute commute for most participants.

2.2 Audit Study: Measuring Anti-favela Discrimination

There is little experimental evidence on whether employers discriminate against favela jobseekers. In Rio, Westphal (2014) conducted an audit study with résumés from different favelas and found no discrimination on average – but with some geographical heterogeneity.[8] Since the Westphal (2014) estimates were ten years old, we conducted a new audit study estimating anti-favela discrimination in callbacks for entry-level sales jobs – similar to the jobs used in our experiments with jobseekers.

We created four fictitious worker profiles with complete high scholl, two male and two female. Age, job experiences, sales-related certifications, and résumé templates varied slightly across profiles. A local consultant with experience matching young favela residents with formal jobs revised these profiles to ensure they were competitive but not unrealistic.

For each profile, we created two copies that differed in name, email, phone number, and address – one from Maré and one from Bonsucesso, which is a non-favela neighborhood adjacent to Maré (see Appendix C for an example). We selected addresses that unambiguously mapped to either Maré or Bonsucesso, keeping the estimated commuting time constant. Maré is a widely recognized favela in Rio, so employers can immediately tell the neighborhood is a favela. Information about the Maré-Bonsucesso callback gap is also relevant for jobseekers in Manguinhos and Jacarezinho, since they update their beliefs about their own neighborhoods similarly to Maré residents when learning about the audit study results (see Figure A.12).

We collected sales job postings (e.g., sales associate, telemarketing salesperson) no older than two weeks from five popular job search websites.[9] Then, research assistants applied to each job posting with two different profiles, with randomized addresses.[10] We submitted 1,400 applications to 700 jobs between February and May 2023. Research assistants monitored the phone numbers and emails until the end of June and coded all non-automatic, non-negative

[8] Zanoni et al. (2023) used the incentivized résumé rating method (Kessler et al., 2019) to measure anti-favela discrimination in Argentina, finding substantial discrimination.

[9] The websites were Catho, Indeed, Infojobs, LinkedIn, and Riovagas. If a posting listed a requirement that one or more of our profiles did not have, or if it was more than two hours away from our addresses by public transport, we also discarded it

[10] The exact randomization procedure was that, for each job posting, we first randomly ordered the four profiles. Among the first two randomly ordered profiles, we randomly picked one for being from Maré. We did the same for the latter two, which were back-ups. A research assistant applied to each posting with two profiles, following the order. The back-up profiles were only used for gendered jobs. If a job were gendered, the research assistant would still follow the suggested order but skip the profiles of mismatched gender. This skipping happened in 9% of postings, and results are similar if we drop those.

replies as callbacks.

The resulting callback rates are similar across neighborhoods: for favela resumes, it is 19.3%, while for non-favela resumes, it is 19.6%, giving a 0.3 p.p. difference between them (p=0.38 to 0.87, depending on the specification, see Table C.1 for details). These similar callback rates do not imply an absence of discrimination against favela residents. For instance, if recruiters believe favela residents are *ceteris paribus* more likely to accept a job offer, that might offset callback differences caused by anti-favela taste-based discrimination (Kessler et al., 2019). Another explanation for the results is that firms are sophisticated and anticipate that some Maré residents obfuscate their neighborhood and instead say they live in Bonsucesso (as we observe in our experiments discussed below), making the declared address uninformative. Nevertheless, even if the audit study measure does not reveal the "real" discrimination level, it provides an objective benchmark for measuring whether jobseekers under- or overestimate anti-favela discrimination.

2.3 Perceived vs. Actual Discrimination

In our door-to-door survey – discussed in detail in the next section – we collected incentivized predictions of what callback rates we would find in our audit study (similar to the method used in Haaland and Roth (2021)). We focus on predictions about the jobseekers' favela of residence versus the adjacent non-favela neighborhood and compare that with the observed Maré and Bonsucesso callback rates.[11]

The top panel in Figure 2 compares callback rate predictions against the ones found in the audit study. On average, jobseekers predict a callback rate of 63% for their adjacent non-favela neighborhood, with 81% predicting callback rates of at least 50%. Jobseekers' guesses are closer to the audit estimates when estimating callback rates for favelas but are, on average, too optimistic: the average prediction for one's favela callback rate is 30% – over 50% larger than the audit study estimates.

The bottom panel in Figure 2 shows the distribution of implied discrimination rates, i.e., the percent drop in callback induced by having a favela instead of a non-favela address. We see that 87% predict discrimination (i.e., a decrease in callback), and 84% predict decreases larger than the upper bound of our 95% confidence interval for the discrimination rate in the audit study. The median jobseeker predicts a 50% discrimination rate, substantially more than the 17.5% upper bound given by our audit study.

We consider jobseekers' predictions of the audit results to be our best measure of expected

[11] We reach similar conclusions if we instead always use beliefs about Maré and Bonsucesso, which are the audit study neighborhoods, see Figure A.10)

anti-favela discrimination, since it is an incentivized and objective benchmark. Reassuringly, providing information on callback rates also decreases an incentivized measure of discrimination regarding the partner HR firm (see Section 4.2), and expected discrimination in the audit study strongly correlates with i) a Likert measure of anti-favela discrimination and ii) a "personalized" measure of discrimination comparing expected future employment probability for oneself against a "clone" of the respondent in the adjacent non-favela (see Figure A.11).

In our survey, we also asked most (N=1,497) jobseekers about the main reasons why employers would discriminate against favela residents. Jobseekers mentioned a mix of productivity- and taste-based reasons. The most common reasons were loss of workdays because of police raids (mentioned by 74%), racism (68%), dislike because of cultural differences (e.g., speech) (66%), and dislike of favela residents (65%). Hence, favela jobseekers have rich second-order beleifs about employer. Notably, jobseekers think that employers understand and act on the correlation between address and other sources of discrimination like race, which might explain the race heterogeneities in our treatment effects.

3 Experiment Design

In early March 2023, before any randomization in the supply-side experiments, we pre-registered the Address Omission Experiment and the Interview Experiment. Those experiments test whether expected discrimination can affect job-seeking behaviors by randomizing whether jobseekers may expect the source of discrimination (address) to be visible or not to the employer. As our treatments typically reduce expected address visibility in relation to the status quo, one can think of these two experiments estimating the effects of "blinding" policies on job application rates and interview performance.

In June 2023, after learning from the audit study that callback rates were similar for favela and non-favela neighborhoods, we decided to launch the Information Experiment. The information from the audit study provided a natural way to randomize market-level expected discrimination (i.e., sharing the audit results) and answer whether providing such information could encourage more people to apply. Since we were concerned about a potentially weak first stage in the Address Omission Experiment at the time, we decided discontinue it in favor of introducing the Information Experiment, which had an easily-verifiable first stage and tested a type of policy that could be transplanted to other contexts. We amended our pre-registration as we introduced this third experiment. This was the single major change in our pre-registered plans (see Appendix B for more details).

We also pre-registered four heterogeneity analyses: by expected discrimination, race, skill, and gender. The heterogeneity by expected discrimination is key to confirming our mechanism of interest. For comparisons, we define the group of jobseekers expecting high discrimination as those who expect 50% discrimination or more when predicting the audit study (i.e., at or above median).[12] The race heterogeneity allows us to observe how correlated stigmas interact. The skill heterogeneity could tell us how expected discrimination changes the talent pool available to employers, and the gender heterogeneity could inform us about whether favela males – who are more likely to be gang members – or females react more to expected discrimination. We discuss heterogeneity by expected discrimination together with average treatment effects (since it is our mechanism of interest). As our findings suggest race plays a major in determining the effects of discrimination, we also discuss those in the main text. The remaining estimates are presented in Appendix A.

3.1 Sample Recruitment

Partners. To advertise real jobs to participants, we partnered with one of Latin America's largest cosmetics franchise and retail chains. This firm was interested in increasing its penetration into favelas and diversity among its workers, allowing us to advertise three entry-level sales jobs. They committed to giving full consideration and fast-tracking promising applicants recruited through our pipeline. We also had the support of several NGOs in each favela. These institutions were extremely important since they had access to local networks and provided feedback on our survey, logistics, and research questions. Crucially NGOs' networks allowed us to hire and train surveyors locally and facilitated obtaining the approval of residents' associations – the relevant political brokers in favelas.

Sampling. Surveyors worked door-to-door to identify favela jobseekers who: i) were between 18 and 40 years old, ii) had completed or were in the last year of high school, and iii) were looking for a *new* full-time formal job. This strategy rules out people who were so wary of anti-favela discrimination and outsiders that they gave up on looking for a formal job or answering an academic survey. So, in a sense, we might be estimating lower bounds on the effects of expected discrimination. To avoid spillovers (since our randomizations are at the individual level) and maximize privacy, surveyors would interview at most one person per household, one-on-one. Every participant received R$5 ($\approx$1 USD) and was entered into a lottery for R$500 (see Figure D.1 for photos of in-progress interviews). Table A.1 presents sample summary statistics: 62% were recruited in Maré, 30% are male, 22% are white, and the average age is 26. In addition, 25% had never worked before, and 32% reported currently working full- or part-time

[12] This definition pools jobseekers who expect fairly high discrimination rates with those who expect none. Nevertheless, results are similar when considering a cut-off of, for instance, 25%.

(most in the informal sector).

Survey. There were four blocks of questions. The first block collected general background information and labor market experience. The second block introduced the HR firm as a partner and asked for the jobseeker's permission to share their basic background with the firm. The third block was about skills. The final block was about anti-favela discrimination and expectations about labor market prospects.

"HR firm". After collecting background information, the surveyor introduced a partner HR firm, which operated in Rio, assisting large companies with their recruitment. The surveyor then asked for permission to share the respondent's answers until that point (i.e., basic demographics and job experience) with the HR firm, so the jobseeker could be invitated to apply for jobs. We (the researchers) operated this HR firm.

Deception. Our choice not to present the HR firm as part of the study was deceptive to the extent that jobseekers could not have anticipated that researchers would observe their interactions with the firm. This was strictly necessary for the design, and the only element of deception in this study. This separation between HR firm and academic researchers served to emulate regular labor market interactions, as research and surveys are commonly linked with local NGOs in favelas. If the research team directly invited respondents to apply for a job, jobseekers could believe they would receive special treatment. At any rate, the HR firm invited jobseekers to apply for real jobs and indeed acted as a recruitment agent.[13]

After choosing whether to share data with the HR firm, surveyors moved to a block on skills. The block started by asking jobseekers whether they had completed courses or training programs relevant to the job market and then asked for self-ratings on computer and soft skills (e.g., punctuality, salesmanship, and leadership). At the end of this block, participants could take an incentivized one-minute test. The test consisted of answering as many basic algebra questions as possible to receive an extra R$0.25 for each correct answer. We use this math test as one of the three components in our skill measure. The other two components are education (self-reported) and communication skills, which are assessed privately by the surveyor on a Likert scale at the end of the survey. We standardize and average these measures to form an index and classify those above the median as "high-skill".

The fourth and final survey block involves questions about job market prospects and expected discrimination. Almost one-third of our sample has heard of somebody who did not get (or lost) a job only because they were from a favela, and a similar number report having

[13] Our debriefing procedures include (i) carefully debriefing those eventually hired by our partner and (ii) inviting participants who applied for the job for a meeting to discuss the study's findings and the use of their data. For the duration of the study, we kept a website and a contact email running, in case any jobseeker searched online for the HR firm.

personally suffered the same. Before initiating the Information Experiment, our survey also included questions on *why* jobseekers believed firms would discriminate against favela residents (mentioned in Section 2.3).

Measuring Expected Discrimination. As our main measure of expected discrimination, we incentivized jobseekers to predict the callback rates we would find for each neighborhood in our audit study (Section 2.2), paying an extra R$100 (≈20 USD) to the ten people who got closer to the true estimates. For both Maré and Manguinhos, we used Bonsucesso as the adjacent non-favela neighborhood. For Jacarezinho, we used Maria da Graça since Bonsucesso is not immediately adjacent (see Table A.2 for Census summary statistics for each neighborhood). As our audit study covered only Maré and Bonsucesso, we elicit incentivized predictions for these other neighborhoods by initially stating that we only knew the correct answer for some of the questions. Figure D.2 shows the complete elicitation script).

Overview. Surveyors completed 2,392 valid interviews, yielding 2,167 eligible participants – 167 did not share their data with the HR firm, and 61 of those who did provided an invalid phone number. Figure 1 shows how the Address Omission, Information, and Interview experiments fit together. We introduced the Information Experiment as we phased out the Address Omission Experiment.[14] As we launched the fieldwork in one favela at a time, the samples for each of the pre-interview studies differ with respect to their favela of origin and some other covariates (see Table A.6 for a comparison). All jobseekers who completed the application form and attended the interview participated in the Interview Experiment.

3.2 Address Omission Experiment (N=1,303)

As the door-to-door survey proceeded, we organized the applicants in batches for the Address Omission Experiment. Every few days, the HR firm sent personalized invitations to apply via WhatsApp to a batch of applicants. Batch sizes varied from 50 to 117 to accommodate logistical capacity. Almost all jobseekers received invitations to apply up to ten days after answering the door-to-door survey.

Treatment. We randomized expected address visibility at the application stage. There were three experimental conditions: *Address Omission*, *Status Quo*, and *Known Address*. Applicants in *Address Omission* received a WhatsApp message from the HR firm inviting them to apply and saying that a home address **is not** needed for applying. Those in *Status Quo* and *Known Address* received a message saying an address **is** needed (see below). The difference between the two conditions in which address is needed is that in *Status Quo* the jobseeker fills in the address (the

[14]There was an overlap of 174 participants between the two pre-interview experiments during the phase-out. For simplicity, the main text presents results for the non-overlapping samples.

common practice in our context), allowing us to observe how often applicants obfuscate their real addresses. In *Known Address*, the form states that the research team has already shared the jobseeker's address (besides name and phone number), so applicants just need to double-check it. Hence, in *Known Address*, we make sure that obfuscation is not possible, allowing us to test whether making one's favela address fully visible affects application behavior (see Figure D.4 for the differences across forms).

WhatsApp Invite Message:

> *Hi [NAME], how are you? This is Vanessa from **SAM HR**. I'm contacting you because you are one of the people in our database who fits the requirements for some of our vacancies. In addition to salary, these jobs offer benefits such as daycare and health insurance.*
>
> *You have been selected to participate in one of our streamlined processes! In this stage, you need to provide your **education and any courses or experiences. Your home address is [NOT/ALSO] required.***
>
> *It takes just **5 minutes**! Personal link: go.samrh.com/lyhW1DS5*

The application form started with a brief description of three full-time jobs: i) (in-store) Sales Consultant, ii) Direct Sales Promoter, and iii) Direct Sales Supervisor (see Figure D.3 for full job descriptions). Then, it asked for the jobseeker's name, phone number, and address – except in *Address Omission*. Then, it proceeded as a standard application form, asking about job experiences, skills, and motivations. Finally, the jobseeker declared their availability for an interview.

Outcomes. Our main pre-registered outcomes are whether the jobseeker completes the online application form and attends the job interview, which typically within two weeks of each other. While not an experimental outcome per se, we also calculate the address obfuscation rate for those in the *Status Quo* arm. We consider that a jobseeker has obfuscated their address if the declared neighborhood is neither a favela nor the postal service neighborhood of the jobseeker's real address (recorded by the surveyor in the door-to-door survey).

Conceptualization. As experimental conditions differed only within the job application procedure, it is reasonable to assume that this treatment only affects the expected value of applying to the jobs in the experiment. We can think of a a jobseeker that assigns value V, success probability p, and has an application cost c, applying if $pV - c > 0$ (normalizing the outside option to zero). Then, the treatment shifts pV, since the differences in the application procedure are minor. For instance, in *Address Omission* perceived pV might be larger both because a jobseeker perceives a higher success probability and because they will be less likely to suffer address-based discrimination on the job.

3.3 Information Experiment (N=690)

The Address Omission Experiment ran until May 2023. As we phased it out, we embedded the Information Experiment in our door-to-door survey. The reasons for introducing this experiment are described at the beginning of Section 3, but here we note another advantage: providing market-level statistics in the survey (detached from the application procedure) sidesteps confounders related to a change in the application process also causing changes in how people perceive the HR firm in ways that are unrelated to expected discrimination. Such confounders are the reason why studies randomizing the language used in jobs ads like Del Carpio and Fujiwara (2023), Burn et al. (2022), or whether AI is used to review applications, like Avery et al. (2023), can not sharply identify the effects of expected discrimination. For instance, in Del Carpio and Fujiwara (2023), gender-neutral language could imply less gender-based discrimination, but it also suggested better work-life balance, which can appeal differently to males and females.

Treatment. We randomized participants into three treatment arms: (i) *No Info*, in which no information was displayed, (ii) *Favela Info*, in which we disclosed only the favela's callback rate (19.3%, from our audit study), and (iii) *Full Info*, in which we showed both the favela and non-favela callback rates (19.3% and 19.6%) – thus revealing that we found no discrimination in callback rates. *Full Info* was designed as our main treatment decreasing expected anti-favela discrimination, but it also reveals the favela callback level. We included the *Favela Info* condition as an alternative control, holding the knowledge of the favela callback level constant. See Figure 3 for the graphs the surveyors used to convey the treatment.

Similar to the Address Omission Experiment, the HR company invited respondents to apply for our partner's jobs. There were only two differences. First, to emulate the most realistic application procedure, we only use *Status Quo* procedures (i.e., we ask applicants to provide their home address). Second, since there was no randomization in the application procedure, we could decrease the batch size and invite jobseekers to apply more often, one to four days after they answered the door-to-door survey.

Endline survey. We conducted an endline survey over WhatsApp to check whether the belief shift caused by the Information Experiment persisted and to collect a self-reported number of job applications sent after answering the door-to-door survey. Participants were contacted two weeks after they answered the door-to-door survey. To maximize responses, participants entered a lottery for R$200 (≈40USD), and we only asked multiple-choice questions.

Outcomes. Besides the application progress outcomes used in the Address Omission Experiment, we also pre-registered as main outcomes the self-reported number of applications sent after two weeks, address obfuscation and immediate belief updates as main outcomes. As updated beliefs, we chose the incentivized predictions of what callback rates the partner HR firm would implement in each neighborhood. There is no ground truth for these callback rates, since we operated the HR firm and invited only favela jobseekers to apply. We incentivized these beliefs by including them in the set of questions in which we elicited beliefs about our audit study callback rates. The surveyor introduced this set of questions with a statement clarifying that we only knew the answer to *some* of the questions.

Conceptualization. Similar to the experimental conditions in the Address Omission Experiment, we may think of the information treatments as shifting the expected callback probability and job value – as jobseekers might value formal job offers more so once they learn that they were too pessimistic about anti-favela discrimination. The main difference is that the information is relevant for *all* jobs. In this case, shifting the expected callback level p can have a non-monotonic effect on application rates. Intuitively, at a low p, an increase in p makes a marginal application much more valuable, so it worth applying to more jobs. But, if you already expect to receive "enough" callbacks, an increase in p allows you to decrease the number of costly applications while still getting enough callbacks.[15] We designed the *Favela Info* condition to fix p, and *Full Info* could be thought of as decreasing expected discrimination (e.g., favela residents now expect to receive impartial treatment during the application process and on the job) and increasing the value of all matches after fixing p. Obfuscation rates should be increasing in expected discrimination and decreasing in own-neighborhood expected callback rate.

3.4 Interview Experiment (N=422)

The HR firm invited all jobseekers who completed the application form for a job interview in an office in Downtown Rio. Attendees received a R\$25 ($\approx$5 USD) transport subsidy – enough to cover transport fares. We rented a reception desk and meeting rooms in a co-working space, so applicants first had to go through the building's reception and then arrive at the co-working floor. Interviews took ten to fifteen minutes each, and we scheduled them with enough of a gap so that jobseekers would rarely, if ever, meet or interact at the premises. See photos in Appendix D.5.

Interview. We hired an experienced HR consultant to revise our interview script and train

[15] To see that, let n be the number of applications to be submitted, p be expected callback probability, c a constant marginal cost and the callback value $V(n,p)$ be such that $V_n > 0$ and $V_{nn} < 0$. If the jobseeker maximizes $V(n,p) - nc$ finding an internal solution, the inverse function theorem yields $\frac{\partial n^*}{\partial p} = -\frac{V_{np}(n^*,p)}{V_{nn}(n^*,p)}$, which has the same sign as $V_{np}(n^*,p)$. Taking, for instance, a jobseeker that only cares about getting the first callback, i.e., $V(n,p) = 1-(1-p)^n$, then one can have $V_{np}(n^*,p) > 0$ for low p and $V_{np}(n^*,p) < 0$ for high p.

our two interviewers. The script contained a set of standard interview questions for sales jobs, including questions about strengths, weaknesses, comparative advantages, past work experiences, and an activity where the applicant had to pick an item and provide a sales pitch for it (see Appendix D.1 for details). Interviewers were instructed to stick to the script and act the same towards all candidates.

Treatment. We randomize expected address visibility at the job interview. A receptionist greeted arriving candidates and asked to confirm their name, date of birth, and address, and told them to wait. Moments later, the receptionist told the jobseeker that the interviewer was ready, and, to keep the process objective, the interviewer "will only know your name" (*Name-Only* condition) or "will only know your name and address" (*Name-and-Address*). Hence, the conditions differed by two words only: "and address". Interviewers were blind to the whole procedure until the end of all interviews, so any effects on the interview must initiate with the candidate. Later, we debriefed the interviewers both to learn their impressions and to avoid participant deception – i.e., the receptionist's statement was ambiguous about when the interviewer would learn about the addresses. Note that our design rules out self-signalling mechanisms (e.g., that when a person is reminded of their address, they lose confidence on their abilities), since all candidates are asked to confirm address before treatment.

Outcomes. The interviewer evaluated candidates immediately after each interview, and interviewees filled out a form with self-assessment questions at the reception desk before receiving the transport subsidy. Interviewers coded, on 0–10 scales, i) how well the interviewee performed overall, ii) how nervous the interviewee was (reverse-coded as calmness), and iii) how professionally the interviewee behaved. Interviewees filled out self-assessments for the same three dimensions. We construct z-scores for each of the six dimensions by normalizing the scores by the mean and standard deviation of those in the *Name-and-Address* condition. For the interviewer-assessed dimensions, we normalize interviewer-wise to account for fixed effects and dispersion differences across interviewers.

To maximize statistical power and reduce the risk of multiple hypothesis testing, we construct an inverse-covariance-weighted index of impressions for the interviewers and for the interviewees (Anderson, 2008). As our primary aggregate performance measure, we average the two. While this averaging mixes impressions of different relevances – i.e., the interviewer's impressions matter for the jobs at hand, while the candidate's impressions matter for their future beliefs, say about whether they should apply to similar jobs again –, it allows us to extract a more accurate signal. We also present broken-down estimates.

Conceptualization. The treatment shifts the candidate's second-order beliefs about how the interviewer might see them. Candidates might respond to that strategically or involuntarily.

For instance, thinking that the interviewer knows one's address might generate feelings of unfairness, leading to automatic stress responses (Berger and Sarnyai, 2015). Such stress, along with the high stakes of a job interview, could lead to choking under pressure, a phenomenon documented in multiple sport competitions (Böheim et al., 2019; Harb-Wu and Krumer, 2019; Teeselink et al., 2020). If a candidate simply believes their performance will be heavily discounted due to discrimination (i.e., lower returns to effort), the optimal response might be to try harder to impress, which can also lead to stress, or to disengage and reduce effort is the barrier is perceived to be insurmountable. An initial effect may compound or dissipate, depending on how interviewers deal with it (e.g., they might "smell blood in the water", further increasing stress).

3.5 Randomization, Balance, and Estimation

Randomization for the Address Omission Experiment proceeded in batches. We stratified by expected discrimination (batch-wise), with equal probability of each treatment within and across strata. We proceeded similarly for the Interview Experiment, randomizing in batches after jobseekers completed the application form.[16] The offline survey app on the surveyors' tablets implemented the randomization for the Information Experiment on the spot – also with equal probabilities. All randomizations were independent across experiments.

Tables A.3, A.4, and A.5 display randomization balance checks. Given the necessity of randomizing batch-wise (for the Address Omission Experiment and Interview Experiment) or on the spot (for the Information Experiment), we could not stratify on multiple variables or at all in the latter case. Hence, we see some imbalances. Out of the 45 comparisons to the "control" groups in tables A.3, A.4, and A.5, one is significant at the 1% level, three at the 5% level, and four at the 10% level, which is not far from what one would expect from randomness. Results with controls are very similar; versions of the main tables and figures using double-lasso selected controls can be found in Appendix A.

To test for the effect of expected stigma visibility in the application procedure and plot the average outcomes of each experimental group, we estimate a saturated model:

$$y_i = \beta_{SQ} Status\ Quo_i + \beta_{KA} Known\ Address_i + \beta_{AO} Address\ Omission_i + \varepsilon_i \tag{1}$$

where $y_i \in \{0, 100\}$ (to yield percentages), and each coefficient captures the outcome level for each treatment group. For the Information Experiment, we use the same specification as in

[16]Due to logistical issues, we had to randomize the treatment status of ten participants as they arrived at the interview office. Results are similar after dropping those ten.

Equation 1 (i.e., one indicator for each treatment). Our interview performance outcomes are normalized z-scores, or their inverse-covariance-weighted averages (Anderson, 2008). Hence, only differences across groups are informative, and we simply regress each outcome on a treatment indicator for *Name-Only* and a constant (i.e., our estimates show the effect of reducing expected discrimination on interview performance). We present robust standard errors for all models, using the HC3 variance-covariance matrix estimator (Long and Ervin, 2000).

4 Results

4.1 *Address Omission Experiment*

Expected address visibility does not affect average job application rates (left panel, Figure 4). If expected address visibility leads to expected discrimination which in turn discourages applications, *Address Omission* should have the highest application rate, and *Known Address* the lowest. Instead, we see little variation across conditions: form completion rates hover from 41% to 45% and interview show-up rates are just below or at 20%. The p-values for tests of equality between any two conditions for application outcomes are all above the conventional significance thresholds. We see a similar pattern even when conditioning on those expecting high discrimination at baseline (right panel in Figure 4), providing no evidence that expected discrimination affects average application rates.

At the same time, we see evidence that jobseekers expected discrimination when applying, and that our treatments shifted expected address visibility to some degree. In *Status Quo*, applicants were free to declare their addresses, and we see 25% of all participants declaring obfuscated addresses (among those who finished the application, that rate is 45%). This high obfuscation share suggests that jobseekers indeed expected some address-based discrimination. At the same time, the *Known Address* treatment was effective in preventing obfuscation, since no jobseekers in that condition even tried to provide a "corrected" address with an obfuscated neighborhood (which they could do in the application form). Hence, we have some evidence that our manipulations worked as intended, shifting expected address visibility and expected discrimination.

We see three potential reasons for the null results. First, the treatments might still have been too weak, as jobseekers might have believed that recruiters would *eventually* figure out their neighborhood of origin, so that any gains from hiding address in the initial stage would be erased. Second, there might have been unintended offsetting effects, e.g., *Known Address* might have led some jobseekers to believe the HR firm would favor favela jobseekers, since they were invited *despite* the firm knowing their addresses. Third, expected anti-favela discrimination might not be marginal in the application decision. For instance, non-white jobseekers might expect discrimination either way, because of their race and because race predicts address.

Out of those three reasons, we see more evidence supporting the third one. We find effects of *Address Omission* only in the white subsample (Section 4.4), which is evidence against a weak treatment and in support of low average elasticity of applications to expected discrimination (since 77% of the sample is non-white). Further, our Information Experiment avoids issues related to the first two reasons, since it shifts beliefs about *market-level* discrimination, and, as we also find an average null effect (see next Section), there is more evidence that shifting expected discrimination does not necessarily shift the average application rate. Finally, two exploratory door-to-door survey questions, introduced as we phased out the Address Omission Experiment, also reinforce the idea that expected anti-favela discrimination is not marginal for application decisions: (i) 70% of respondents (N=670) agree that one should apply for all possible postings to do well in the labor market, and (ii) 80% agree that to do well, one should not ruminate about employer discrimination. Hence, many jobseekers seem to try keep expected discrimination out of their minds when applying to jobs.

4.2 Information Experiment

In this experiment, we find that the information treatments successfully shift expected discrimination in callbacks by "correcting" the beliefs about the callback rate for favela and non-favela neighborhood. Nevertheless, we do not find an effect on average application rates.

The information treatments successfully "corrects" beliefs about the callback rate that the HR would implement in favelas: those who initially under- and over-estimate callback converge in their posterior beliefs. Learning *Favela Info*, i.e., that the callback rate for the favela was 19.3%, does not change the *average* expected callback rate for jobseekers' own neighborhoods (Figure 5), but that is due to the belief convergence. For instance, considering only those who initially overestimate the favela callback rate, the average expected callback rate goes from 41% in *No Info* to 37% in *Favela Info* (p=0.09, see Figure A.4 for effects on under- and overestimators of the favela callback rate). When jobseekers learn both callback rates in *Full Info*, underestimators become even more optimistic about their own favela callback rates, and overestimators become more pessimistic. For both subgroups, there is a statistically significant shift in expected callback for one's own favela when learning *Full Info*. Hence, jobseekers use favela *and* non-favela information to update about favela callback rates.

Considering beliefs about the adjacent non-favela callback rate, we also see that jobseekers use information on both favela and non-favela callback rates to update. Since 92% of the sample overestimate the non-favela callback rate, that update is evident even when looking at the full sample in the top-right of Figure 5. Hence, both *Favela Info* and *Full Info* decrease expected discrimination, and the decrease is larger for *Full Info* since it provides more information. The average posterior discrimination rate for the *No Info*, *Favela Info*, and *Full Info* groups are, respectively, 35%, 28%, and 15%, with group differences significant at the 5% or 1% level. The top-right graph in Figure 5 shows a similar pattern for the subsample who expected high discrimination at baseline.

Nevertheless, we do not find statistically significant differences in application rates across information conditions, even for the high expected discrimination group (Figure 5, bottom row). In the same figure, considering obfuscation rates, we see the hypothesized pattern, (i.e., highest obfuscation in the *Favela Info* condition, when most people learn they were too optimistic about their neighborhood's callback rate). That said, the only (marginally) statistically significant difference is between *No Info* and *Favela Info*, conditional on expecting high discrimination. In that case, those receiving *Favela Info* obfuscate 13% of the time, more than double the share in *No Info* (p=0.1).[17]

Our endline survey generally confirms the findings above. There was no differential attrition in participation (column (1) in Table 1), so sample selection into the endline should not be an issue. There is evidence that the decrease in expected discrimination caused by *Full Info* persists for at least two weeks, at least in comparison with *Favela Info* (p=0.06, see column (2) in Table 1)). In a pooled comparison of Full Info against the two other arms (not shown in the table), we see p=0.09. Nevertheless, we still see null effects on application rates, but now on a self-report of the total number of jobs the respondent applied for in the last two weeks.

4.3 Interview Experiment

Even if expected discrimination does not affect average application rates, it could still damage interview performance, since there are many differences between making application decisions and controlling behavior in interviews. For instance, during the interview, the jobseeker must quickly adjust behavior in response to the interviewer, who directly observes and judges performance, making the interview interaction very different from the "cold" decision of whether to apply.

Decreasing expected address visibility, i.e., hearing that the interviewer will only know the

[17] If we assume that the only way in which our treatments affect applications is through beliefs about callback rates, we can estimate the effect of those beliefs with two-stage least squares. Taking that approach lets us exploit variation in how both information treatments affect callback beliefs for favela and non-favela neighborhoods, but leads to similar takeaways (see Table A.10).

candidate's name, increases the average aggregate performance index by 0.13SD (Figure 6, first estimate). That index averages the interviewer's and the candidate's opinions, and when we break it up by those two components we see that the increase in the self-assessed index (0.17SD) is statistically significant at the 1% level while the increase in interviewers' assessments (0.09SD) is not. While this suggests that the effects on self-assessments are larger, we cannot reject that they are equal (p=0.33). Hence, we have evidence that expected discrimination is on average detrimental to candidate's self-assessments, which may determine their beliefs about getting this job or a similar one in the future, but we are less confident about effects on average interviewers' opinions, which matter directly for hiring decisions. Notably, the average effects of reducing address visibility appear beneficial across the board, as estimates of regarding each of the six index components, although noisy, always go in the same direction (Figure 6, gray circles).

Breaking up the estimates by the groups who expect higher or lower anti-favela discrimination at baseline gives stronger evidence for a self-fulfilling prophecy powered by expected discrimination. For the group expecting high discrimination at baseline we estimate statistically significant increases in performance of over 0.2SD in response to *Name-Only*, in both the candidate's and the interviewer's points-of-view (red diamonds in Figure 6, $p<0.05$ for both index outcomes). While these effects are large for a treatment changing only two words said by the receptionist, they are not unreasonable in magnitude, since having some college education is correlated with a 0.55SD increase in performance. For the group expecting lower discrimination to begin with, we cannot reject the null of no effects at the 5% level for any index or their components (orange squares, Figure 6). Hence expecting to have a visible address leads to worse performance among those who expect high discrimination from the start, even when interviewers do not have the information to discriminate. Further, the size of the differential effect on people who expect high discrimination is stable when we allow for other heterogeneity dimensions (Table A.8), suggesting that other characteristics correlated with expected discrimination are not responsible for the observed effect heterogeneity.[18]

Expected discrimination also seems to hurt interviewer-assessed performance at the right tail of the performance distribution among those who expect high discrimination, so it can generate gaps in the composition of hires.[19] Specifically, we can show that among those who expect high discrimination, there is a drop in the share of candidates above different cut-offs the performance distribution. For instance, take as a benchmark the cut-off corresponding to

[18] Results in Figure 6 are similar when including double-lasso selected controls, see Figure A.3.

[19] The ideal exercise here would use data on who the cosmetics firm actually hired. Nevertheless, we only received sparse and incomplete information on which candidates were further contacted and hired, making this exercise infeasible.

the 90th percentile of interviewer-assessed performance index among those in *Name-Only* (i.e., mimicking a world with low expected discrimination). Then, among candidates expecting high discrimination, there are only half as many above that cut-off in *Name-and-Addres* than in *Name-Only* (p=0.09, see Figure 7 for the empirical CDFs). Put another way, among those who expect high discrimination, *Name-and-Address* ejects about half of the candidates from the top 10% of the distribution. At the top 1%, there there are only candidates in the *Name-Only* arm.

While the Interview Experiment was not designed to identify the mechanisms through which expected discrimination affects interview performance, three empirical patterns suggest that a mix of stress and heightened stakes lead candidates to choke under pressure and do worse when their addresses are visible. First, since there is little pressure when a jobseeker is deciding whether to apply or not, choking under pressure is consistent with our observation of effects being more pronounced in the interview rather than at the application decision.

The second empirical pattern consistent with a role for choking under pressure is how the heterogeneous effects by expected discrimination level are distributed across the six index components. We see that the effect of *Name-Only* on self-assessed calmness among those who expect high discrimination is 0.4SD larger than the effect on those who expect lower discrimination at baseline (p=0.04), while there is no effect heterogeneity in the other two types of self-assessed perceptions, and, from the point of view of the interviewer, effect sizes are larger in the high expected discrimination group for all three components. This suggests that increasing expected discrimination or unfairness leads to a stress response (consistent with a wider literature on discrimination or injustice leading to stress, e.g., Berger and Sarnyai (2015)), which then reflects badly on how the interviewer assesses the candidate.[20] While not ruled out, a mechanism in which discrimination reduces the perceived returns to effort during interview would not by itself generate such increases in nervousness when address is visible in the high-expected discrimination group, since address visibility would lower stakes significantly.

Third, we see some evidence that jobseekers might find it difficult to be strategic at the office, which would be consistent with stress making it hard for candidates to perform. For instance, if we look at jobseekers who went through a *Status Quo* application process (manually filling their addresses in the online form) and made it to the interview, we see that the same jobseekers are 20% (5.7 p.p.) less likely to obfuscate their addresses at the interview office (p<0.01). Also, we see that the index component that is most affected (on average) by *Name-Only* is the self-perceived measure of professional behavior, suggesting that jobseekers do not

[20]Note that the effect on nervousness could have gone in the opposite direction if *Name-Only* led to a higher cognitive load among those who try to pass as non-favela residents (e.g., the pressure to be careful and not reveal any hints you are actually from the favela could overpower the nervousness induced by expecting discrimination). Hence, the effect of expected discrimination overpowers this other stress-inducing mechanism.

self-regulate their behaviors as much when they believe their address stigma is visible. While a mechanism involving candidates trying to overcompensate for their address visibility could increase nervousness, it would not explain why we see candidates acting less professional when their addresses are visible.

4.4 Race and Stigma Visibility

Race and favela residence are correlated stigmas, and jobseekers understand that. In Rio, address is highly predictive of race, as white people are a majority of the population outside favelas but only one-third of the favela population. As such, an employer who discriminates racially could use address information to rule out potentially non-white candidates. Favela jobseekers anticipate this employer behavior, as 68% of our survey respondents mentioned racism as an important reason why employers discriminate against favela residents. Then, it is reasonable to expect that manipulating expected address discrimination would affect white and non-white jobseekers differently. For instance, reducing expected address visibility might encourage white jobseekers more than non-white jobseekers, since a white jobseekers of unknown address is much less likely to reside in a favela than a non-white one, and the latter may expect racial discrimination either way.

In Table 2, we present a summary of how the treatments reducing expected discrimination differentially affected white and non-white jobseekers across the experiments. First, note that being told that the interviewer would only know the candidate's name increases the aggregate performance index only for white jobseekers (see column (1), p=0.09 when testing the difference in the effects). This is consistent with the idea of white jobseekers benefiting from passing as non-favela residents, and non-white jobseekers expecting discrimination either way, as mentioned above.

Next in Table 2, we show the effects of removing the need to declare address in the Address Omission Experiment – against the other two conditions pooled, to increase power and simplify interpretation. The most relevant pattern is the statistically significant difference in how removing the need to declare address affects white and non-white jobseekers' application rates (p=0.01) and show-up rates (p=0.05) – see columns (2) and (3). Consistent with the story above, reduced expected address visibility encourages white jobseekers relatively more. While there is a statistically significant *negative* effect of reducing expected address visibility on the application rate for non-white jobseekers, that negative effect fades when we consider the interview attendance outcome (which is more economically relevant). As the point estimates of the effects point in different directions for the white and non-white subsamples (and as the white subsample is relatively smaller), we end up the the average null effects discussed in Section 4.1.

Finally, in the last two columns in Table 2, we show heterogeneous effects of telling people that we found no market-level address-based discrimination in callbacks – again, comparing against the two other conditions pooled to improve power. Here the pattern flips: white jobseekers become relatively discouraged, applying and showing up for interview relatively less (p=0.11 and p=0.02) when they are told about the negligible level of discrimination in callbacks. A natural interpretation for non-white jobseekers becoming more encouraged could be that observing no address-based discrimination also leads to decreasing expected *racial* discrimination (because it is evidence that firms do not use address, which strongly predicts race, as a filtering mechanism). For white jobseekers, the audit study results imply lower returns for passing as a non-favela resident and increased expected competition from non-white applicants. Considering interview show-up, in column (5), we see evidence of strong discouragement among white jobseekers, but not much encouragement among non-white jobseekers.

5 Discussion

Relevance of the effects on self-assessment. Even if the effects of stigma visibility were restricted to self-assessment, they may still have important implications. For instance, after a negative interview experience, jobseekers might be reticent to apply for other jobs that require formal interviews. Also, note that in "regular" interviews, discriminatory behavior among interviewers can exacerbate any effects of expected discrimination – in our interviews, that channel was shut down since interviewers knew only names and stayed on script. Finally, even if we disregard discouragement effects or interview performance completely, expected discrimination can undermine jobseeker's psychological welfare (Pascoe and Smart Richman, 2009; Schmitt et al., 2014), as we show that it leads to negative interview experiences.

Policy Implications. Our experiments have implications for policies that restrict the information recruiters may access. First, consider policies that reduce stigma visibility at the callback stage, such as résumé anonymization, or forbidding employers from requesting some specific information. Our results suggest we should not expect such policies to change applicant behavior substantially or across the board. Our analysis of the interaction between race and address visibility suggests that such policies might only encourage applications for groups who can keep on hiding their stigmas later on, as was the case with white jobseekers in our sample. Since there is also evidence that such procedures can backfire when they lead recruiters to make decisions with incomplete information (e.g., Behaghel et al. 2015; Doleac and Hansen 2020), our results suggest these policies should be treated with even more caution.

On the other hand, there is reason to become more optimistic about "blind" interviews (as in Goldin and Rouse 2000), since we show evidence that simply expecting a blind procedure can improve performance. Our study highlights the importance of jobseekers' *second-order* beliefs, rather than whatever other damage discriminating interviewers may impose. Hence, employers should make sure that jobseekers are fully aware of blinding policies. Furthermore, even if a policy hides one stigma, it may fail to have an effect because another stigma may act as a substitute – as we show in Section 4.4. Hence, policies that hide all stigmas during interviews (e.g., audio-only, text, or metaverse interview rounds) could dominate alternatives. AI-intermediated candidate selection is also a promising alternative, as shown in Avery et al. (2023).

If choking under pressure, stress, and difficulty in managing behavior are to blame for the observed effects on interview performance, there is also a case to be made for interventions helping jobseekers to signal their skills more confidently. Abebe et al. (2021) and Carranza et al. (2020) show that giving jobseekers an easy way to show their skills (e.g., with an informative letter of recommendation) has positive effects on employment. Having access to such signalling mechanisms may reduce the weight put on interview performance, and give jobseekers more confidence. Interventions focused on controlling behavior and decreasing anxiety, such as cognitive-behavioral therapy, could also counter such negative effects in interviews.

Intersectionality. Our results show that interventions ignoring race as a correlated source of discrimination can lead to heterogeneous and unintended results, such as increasing racial inequalities. This resonates with the idea that overlapping sources of discrimination compound in ways that can not be summarized by simple additive effects, and that first-best policies are personalized (Carvalho et al., 2022; Crenshaw, 1989). Nevertheless, our findings also suggest predicting this intersectionality (the more personalized treatment effects) is hard. For instance, a policymaker could expect that hiding address at the interview would benefit non-white more than white jobseekers, which would happen if jobseekers thought they could overcome one but not two "strikes" against them – but we find the opposite. Another implication of our findings is that the status quo policy of asking for address at the application stage is damaging only for white jobseekers, which might go against many's intuitions. Hence, an information-constrained policy maker could justifiably pick race-blind policies to address anti-favela expected discrimination as a first-order approximation, and then iterate based on its results.

Implications for firms. Firms may also play a role in decreasing expected discrimination and creating an environment where they can extract a better signal from interviews. For instance, making the candidate-selection process more transparent and credibly committing to non-discriminatory practices (such as diversity, equity, and inclusion). While firms need to consider the trade-offs involved in adopting these policies, our evidence on interview performance suggests that such policies may help firms hire the best talent.

6 Conclusion

This paper provides evidence that expected discrimination can exacerbate the impacts of employer discrimination on jobseekers and work as a self-fulfilling prophecy in job interviews, potentially contributing to labor market inequalities observed in administrative data. It shows that there are still more ways in which discrimination can act as a self-fulfilling prophecy, besides the already-documented channels involving human capital acquisition and on-the-job discrimination. Also, our findings that different policies tackling expected anti-favela discrimination generate race-heterogeneous effects suggest that policymakers should be mindful of overlapping perceived disadvantages, while keeping in mind that correlated sources of discrimination can interact in counterintuitive ways.

Given the relevance of the topic for firms and policymakers, we see an avenue for future research aiming to understand precisely why expected discrimination is (more) relevant at the interview stage. Our results suggest that choking under pressure might be behind the negative effects on interview performance, and experiments varying pressure, or whether the interview has a face-to-face element, could shed light on mechanisms. Moreover, since many institutions have become committed to diversity, equity, and inclusion (DEI) in recent years (Pew Research, 2021; Fath, 2023), an immediate question is whether making such public commitments can indeed decrease jobseekers' expected discrimination regarding those firms. These DEI commitments can be costly for firms (e.g., a firm might need to hire staff to develop and implement such policies), while their upsides are uncertain. If DEI commitments remove a handicap faced by jobseekers who anticipate discrimination and help recruiters in talent identification, they could become more attractive to a broader range of firms.

References

Abebe, Girum, A Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, and Simon Quinn, "Anonymity or distance? Job search and labour market exclusion in a growing African city," *The Review of Economic Studies*, 2021, *88* (3), 1279–1310.

Agüero, Jorge M, Francisco Galarza, and Gustavo Yamada, "(Incorrect) Perceived Returns and Strategic Behavior among Talented Low-Income College Graduates," in "AEA Papers and Proceedings," Vol. 113 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2023, pp. 423–426.

Aksoy, Billur, Ian Chadd, and Boon Han Koh, "Sexual identity, gender, and anticipated discrimination in prosocial behavior," *European Economic Review*, 2023, *154*, 104427.

Anderson, Michael L, "Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects," *Journal of the American statistical Association*, 2008, *103* (484), 1481–1495.

Avery, Mallory, Andreas Leibbrandt, and Joseph Vecci, "Does Artificial Intelligence Help or Hurt Gender Diversity? Evidence from Two Field Experiments on Recruitment in Tech," *Evidence from Two Field Experiments on Recruitment in Tech (February 14, 2023)*, 2023.

Barnes, Nicholas, "The logic of criminal territorial control: military intervention in Rio de Janeiro," *Comparative Political Studies*, 2022, *55* (5), 789–831.

Behaghel, Luc, Bruno Crépon, and Thomas Le Barbanchon, "Unintended effects of anonymous resumes," *American Economic Journal: Applied Economics*, 2015, *7* (3), 1–27.

Berger, Maximus and Zoltán Sarnyai, ""More than skin deep": stress neurobiology and mental health consequences of racial discrimination," *Stress*, 2015, *18* (1), 1–10.

Böheim, René, Dominik Grübl, and Mario Lackner, "Choking under pressure–Evidence of the causal effect of audience size on performance," *Journal of Economic Behavior & Organization*, 2019, *168*, 76–93.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer, "Stereotypes," *The Quarterly Journal of Economics*, 2016, *131* (4), 1753–1794.

Burn, Ian, Daniel Firoozi, Daniel Ladd, and David Neumark, "Help Really Wanted? The Impact of Age Stereotypes in Job Ads on Applications from Older Workers," Technical Report, National Bureau of Economic Research 2022.

_ , _ , _ , and _ , "Age Discrimination and Age Stereotypes in Job Ads," *FRBSF Economic Letter*, 2023, *2023* (07), 1–5.

Carpio, Lucia Del and Thomas Fujiwara, "Do Gender-Neutral Job Ads Promote Diversity? Experimental Evidence from Latin America's Tech Sector," Technical Report, National Bureau of Economic Research 2023.

Carranza, Eliana, Robert Garlick, Kate Orkin, and Neil Rankin, "“Job Search and Hiring With Two-Sided Limited Information About Workseekers’ Skills,” 2020.

Carvalho, Jean-Paul, Bary Pradelski, and Cole Williams, “Affirmative action with multidimensional identities,” *Available at SSRN 4070930*, 2022.

Charness, Gary, Ramón Cobo-Reyes, Simone Meraglia, and Ángela Sánchez, “Anticipated discrimination, choices, and performance: Experimental evidence,” *European Economic Review*, 2020, *127*, 103473.

Coate, Stephen and Glenn C Loury, “Will affirmative-action policies eliminate negative stereotypes?,” *The American Economic Review*, 1993, pp. 1220–1240.

Crenshaw, Kimberlé, “Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics,” *University of Chicago Legal Forum*, 1989.

Dickerson, Andy, Anita Ratcliffe, Bertha Rohenkohl, and Nicolas Van de Sijpe, “Anticipated labour market discrimination and educational achievement,” *The Sheffield Economic Research Paper Series (SERPS)*, 2022, *2022017* (2022017).

Doleac, Jennifer L and Benjamin Hansen, “The unintended consequences of “ban the box”: Statistical discrimination and employment outcomes when criminal histories are hidden,” *Journal of Labor Economics*, 2020, *38* (2), 321–374.

Fath, Sean, “When Blind Hiring Advances DEI – and When It Doesn’t,” *Harvard Business Review*, 2023.

Fryer, Roland G, Jacob K Goeree, and Charles A Holt, “Experience-based discrimination: Classroom games,” *The Journal of Economic Education*, 2005, *36* (2), 160–170.

Glover, Dylan, Amanda Pallais, and William Pariente, “Discrimination as a self-fulfilling prophecy: Evidence from French grocery stores,” *The Quarterly Journal of Economics*, 2017, *132* (3), 1219–1260.

Goldin, Claudia and Cecilia Rouse, “Orchestrating impartiality: The impact of “blind” auditions on female musicians,” *American economic review*, 2000, *90* (4), 715–741.

Haaland, Ingar and Christopher Roth, “Beliefs about racial discrimination and support for pro-black policies,” *The Review of Economics and Statistics*, 2021, pp. 1–38.

Harb-Wu, Ken and Alex Krumer, "Choking under pressure in front of a supportive audience: Evidence from professional biathlon," *Journal of Economic Behavior & Organization*, 2019, *166*, 246–262.

Hoff, Karla and Priyanka Pandey, "Discrimination, social identity, and durable inequalities," *American economic review*, 2006, *96* (2), 206–211.

Kang, Sonia K, Katherine A DeCelles, András Tilcsik, and Sora Jun, "Whitened résumés: Race and self-presentation in the labor market," *Administrative Science Quarterly*, 2016, *61* (3), 469–502.

Kessler, Judd B, Corinne Low, and Colin D Sullivan, "Incentivized resume rating: Eliciting employer preferences without deception," *American Economic Review*, 2019, *109* (11), 3713–44.

Kuhn, Peter and Kailing Shen, "What Happens When Employers Can No Longer Discriminate in Job Ads?," *American Economic Review*, 2023.

Lang, Kevin and Michael Manove, "Education and labor market discrimination," *American Economic Review*, 2011, *101* (4), 1467–1496.

Lepage, Louis-Pierre, Xiaomeng Li, and Basit Zafar, "Anticipated Gender Discrimination and Grade Disclosure," Technical Report, National Bureau of Economic Research 2022.

Lessing, Benjamin, "Conceptualizing criminal governance," *Perspectives on politics*, 2021, *19* (3), 854–873.

Liu, Songqi, Pei Liu, Mo Wang, and Baoshan Zhang, "Effectiveness of stereotype threat interventions: A meta-analytic review.," *Journal of Applied Psychology*, 2021, *106* (6), 921.

Long, J Scott and Laurie H Ervin, "Using heteroscedasticity consistent standard errors in the linear regression model," *The American Statistician*, 2000, *54* (3), 217–224.

Lundberg, Shelly J and Richard Startz, "Private discrimination and social intervention in competitive labor market," *The American economic review*, 1983, *73* (3), 340–347.

Monteiro, Joana, Eduardo Fagundes, Mariana Carvalho, and Ramon Chaves Gomes, "Territorial Criminal Enterprises: Evidence from Rio de Janeiro," Technical Report 2022.

Neumark, David, "Experimental Research on Labor Market Discrimination.," *Journal of Economic Literature*, 2018, *56* (3), 799–866.

Pager, Devah and David S Pedulla, "Race, self-selection, and the job search process," *American Journal of Sociology*, 2015, *120* (4), 1005–1054.

Pascoe, Elizabeth A and Laura Smart Richman, "Perceived discrimination and health: a meta-analytic review.," *Psychological bulletin*, 2009, *135* (4), 531.

Quillian, Lincoln, John J Lee, and Mariana Oliver, "Evidence from field experiments in hiring shows substantial additional racial discrimination after the callback," *Social Forces*, 2020, *99* (2), 732–759.

REDES, DA MARÉ, "Censo de Empreendimentos Econômicos da Maré," *Rio de Janeiro: Observatório de Favelas*, 2014.

Research, Center Pew, "Diversity, Equity and Inclusion in the Workplace," Technical Report 2021.

Riach, Peter A and Judith Rich, "Field experiments of discrimination in the market place," *The economic journal*, 2002, *112* (483), F480–F518.

Rich, Judith, "What do field experiments of discrimination in markets tell us? A meta analysis of studies conducted since 2000," Technical Report 2014.

Ruebeck, Hannah, "Perceived discrimination at work," 2024.

Schmitt, Michael T, Nyla R Branscombe, Tom Postmes, and Amber Garcia, "The consequences of perceived discrimination for psychological well-being: a meta-analytic review.," *Psychological bulletin*, 2014, *140* (4), 921.

Spencer, Steven J, Christine Logel, and Paul G Davies, "Stereotype threat," *Annual review of psychology*, 2016, *67*, 415–437.

Steele, Claude M and Joshua Aronson, "Stereotype Threat and the Intellectual Test Performance of African Americans.," *Journal of Personality and Social Psychology*, 1995, *69* (5), 797.

Teeselink, Bouke Klein, Rogier JD Potter van Loon, Martijn J van den Assem, and Dennie van Dolder, "Incentives, performance and choking in darts," *Journal of Economic Behavior & Organization*, 2020, *169*, 38–52.

UN, Habitat, "Slum Almanac 2015-2016: Tracking Improvement in the Lives of Slum Dwellers.," *Participatory Slum Upgrading Programme*, 2016.

Westphal, Eric, "Urban Slums, Pacification, and Discrimination: A Field Experiment in Rio de Janeiro's Labor Market." Bachelor's thesis, Harvard University 2014.

Word, Carl O, Mark P Zanna, and Joel Cooper, "The nonverbal mediation of self-fulfilling prophecies in interracial interaction," *Journal of experimental social psychology*, 1974, *10* (2), 109–120.

Zanoni, Wladimir, Paloma Acevedo, Giulia Zane, and Hugo Hernández, "Discrimination Against Workers From Slums: What Is its Extent, What Explains It, and How Do We Tackle It?," 2023.

Table 1: Information Does Not Affect Application Rates at Endline

	(1) Responded to endline (0/1)	(2) Exp. discrimination (categorical, 1-4)	(3) # of sent apps (categorical, 1-4)
Favela Info	0.02	0.06	-0.02
	(0.05)	(0.10)	(0.13)
Full Info	0.02	-0.12	0.02
	(0.05)	(0.10)	(0.14)
Observations	690	389	389
Controls	No	No	No
No Info Mean	0.6	2.3	2.5
Favela=Full p	0.96	0.06	0.76

Note: Information Experiment treatment effects on endline survey outcomes. The outcome in column (1) is a dummy for responding the endline survey. The outcome in column (2) takes values from one to four, coding for believing that a favela jobseeker would [NOT suffer=1/suffer A BIT more=2/ suffer A LOT more=3/suffer EXTREMELY more=4] discrimination than someone from the adjacent non-favela when applying to jobs. The outcome in column (3) equals 1 if the jobseeker applied for zero jobs, 2 if applied for a single job, 3 if applied from two to five, and 4 if applied for more jobs than that over the last two weeks. Robust standard errors are shown in parentheses.

Table 2: Race Moderates the Effects of Varying Expected Address Discrimination

	(1) Agg. Performance Index (SDs)	(2) Applied (%)	(3) Showed Up (%)	(4) Applied (%)	(5) Showed Up (%)
$\downarrow \mathbb{E}$[Disc.] $\times$ white	0.31**	9.94	10.14*	-9.37	-15.26**
	(0.13)	(6.05)	(5.20)	(8.71)	(6.25)
$\downarrow \mathbb{E}$[Disc.] $\times$ non-white	0.07	-7.21**	-1.16	6.43	1.90
	(0.07)	(3.30)	(2.56)	(4.57)	(3.73)
Observations	422	1302	1302	690	690
Treatment	$\downarrow \mathbb{E}$[address visibility]	$\downarrow \mathbb{E}$[address visibility]	$\downarrow \mathbb{E}$[address visibility]	$\downarrow \mathbb{E}$[mkt-level disc.]	$\downarrow \mathbb{E}$[mkt-level disc.]
Control Average	-0.0	44.1	18.2	42.3	20.7
White=non-white p-value	0.09	0.01	0.05	0.11	0.02

Note: $\downarrow \mathbb{E}[Disc.]$ stands for a dummy for the experimental condition that in principle should reduce expected discrimination the most in each experiment. That is, $\downarrow \mathbb{E}[Disc.]$ is *Name-Only* in column (1), *Address Omission* in columns (2) and (3), and *Full Info* in columns (4) and (5).

Figure 1: Experimental Design

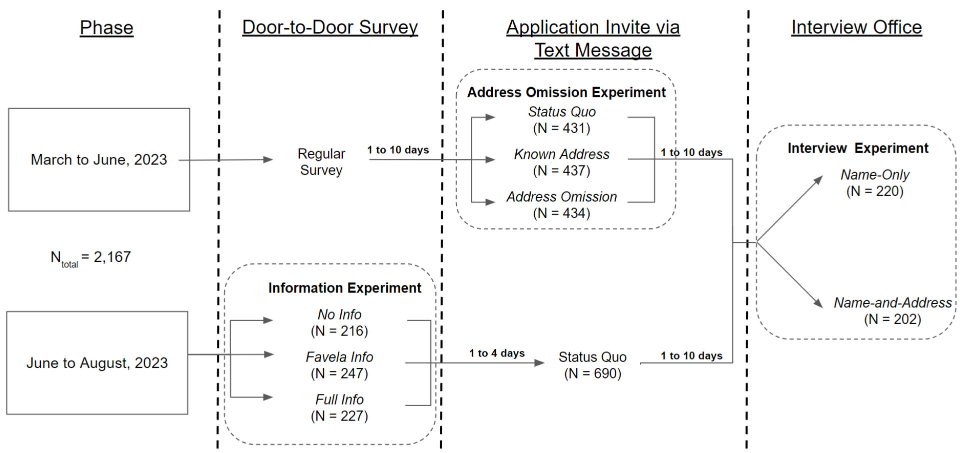

Note: The figure shows a simplified diagram of the flow of participants from door-to-door survey to job interview, for the earlier and later fieldwork periods. See Section 3.1 for details.

Figure 2: Predicted *vs.* Actual Discrimination Rates

Note: The top panel shows the distribution of the guesses for the callback rates in an audit study using résumés with addresses from the respondent's favela or with that favela's adjacent neighborhood. The bottom panel plots the distribution of the implied discrimination rates, measured as the percent drop in callback rate caused by using a favela address. Predictions of more than 50% negative discrimination (i.e., discrimination against non-favela residents) are bunched at the leftmost bin. Vertical dashed lines show the audit study point-estimates. In the bottom graph, guesses are color-coded by whether they fall into the 95% confidence interval of discrimination against Maré (vs. Bonsucesso) résumés (calculated using our audit study).

Figure 3: Information Treatment Delivery

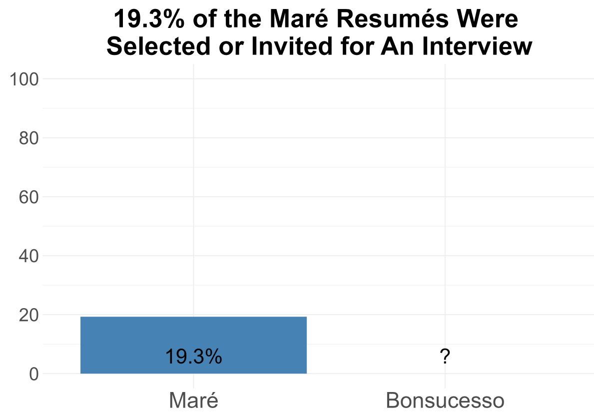

(a) *Favela Info*

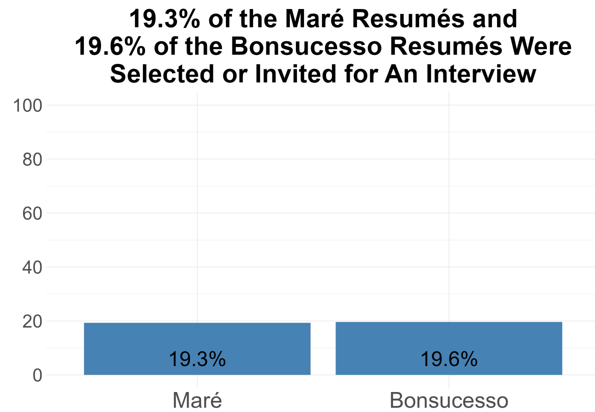

(b) *Full Info*

Note: This Figure shows the images we used to convey the Information Experiment. We showed either one of the plots (or none) to participants immediately after the belief elicitation presented in Figure D.2. The surveyor read the text above each graph when showing it to the respondent.

Figure 4: Address Omission Experiment: No Differences in Application Rates Across Arms

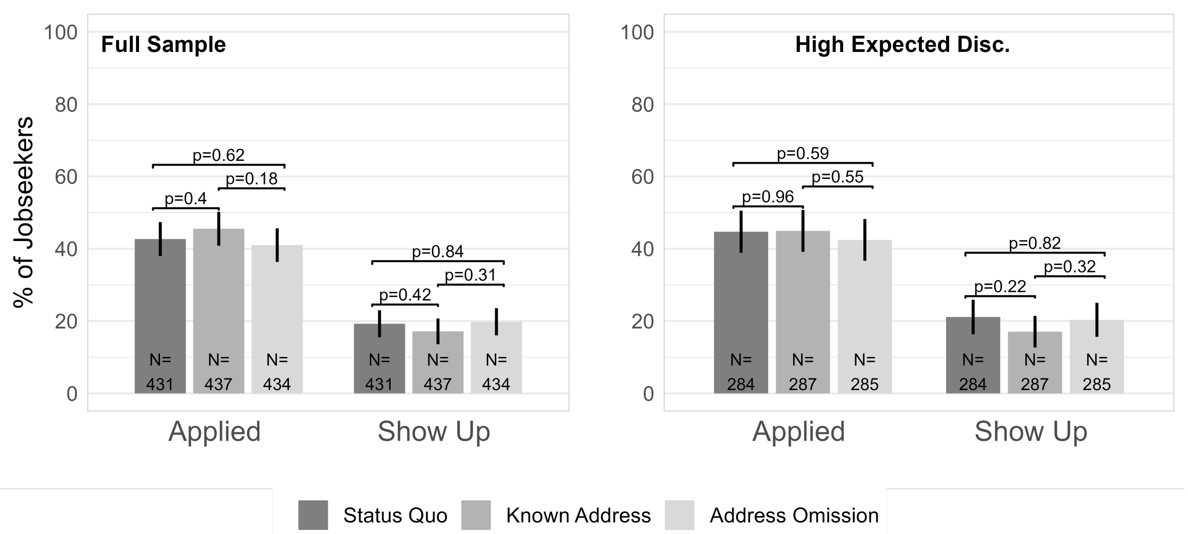

Note: This figure displays shares of all jobseekers in the Address Omission Experiment reaching each stage of the application process. Clicked, means clicking the link in the WhatsApp invite. Applied stage means finishing the online application form, and Show Up means showing up at the interview. The left panel shows results for the full sample, and the right panel shows results conditional on expecting 50% (median) discrimination or more when predicting the audit study. Sample size in each arm is shown at the bottom of each bar. Vertical error bars display 95% confidence intervals, and horizontal bars with tips show p-values for pairwise comparisons above them.

Figure 5: Information Treatment Shifts Beliefs, But Not Interview Show-up

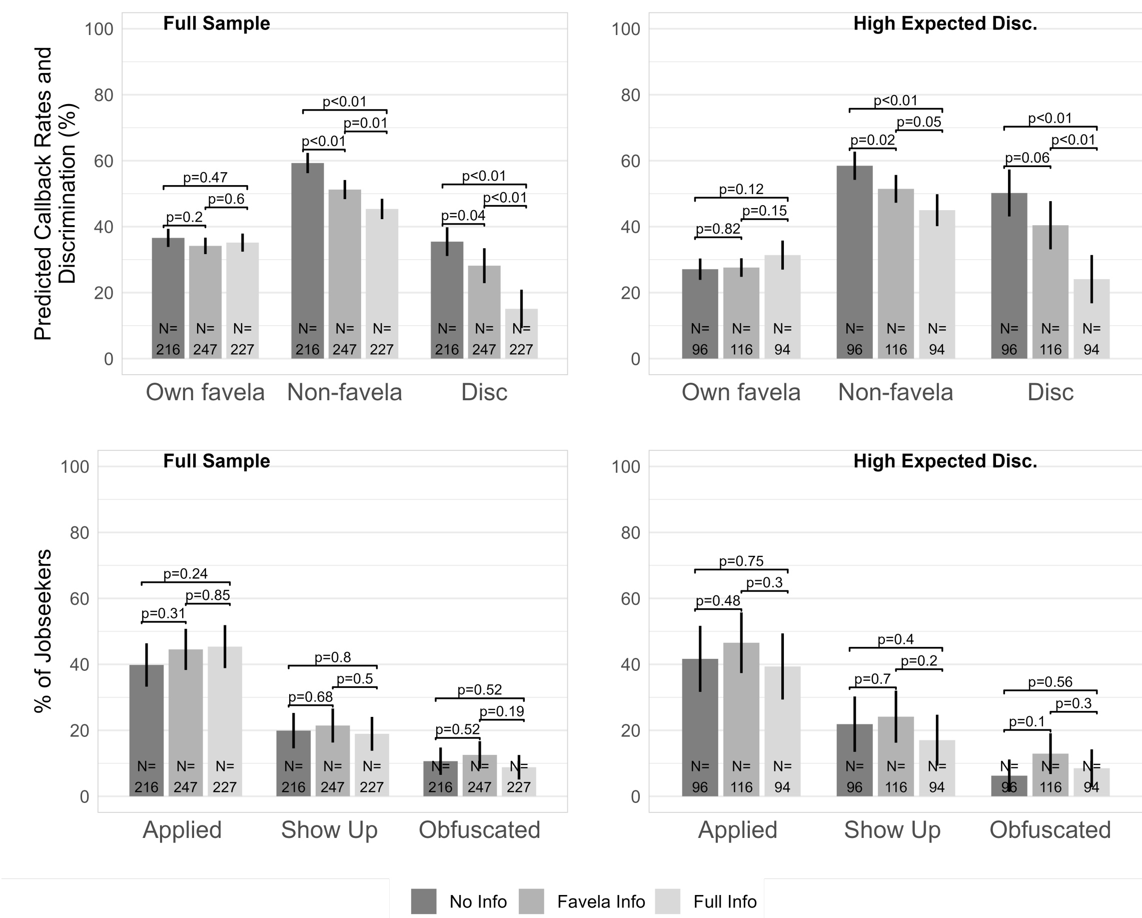

Note: The top row of graphs displays average posterior beliefs of what callback rates the HR firm would implement for jobseekers in each experimental condition. Non-favela and Own favela stands for the callback rate prediction for a respondent's favela and adjacent non-favela. Disc is the implied discrimination rate. The bottom row displays outcomes from the application process. Clicked, means clicking the link in the WhatsApp invite. Applied means finishing the online application form, and Show Up means attending the interview. Obfuscates in app means declaring (in the application form) a neighborhood that is neither a favela nor the postal service neighborhood of the true address. The left column of graphs shows results for the full sample, and the right column shows results conditional on expecting 50% (median) discrimination or more when predicting the audit study. Sample size in each arm is shown at the bottom of each bar. Vertical error bars display 95% confidence intervals, and horizontal bars with tips show p-values for pairwise comparisons above them.

Figure 6: Expected Stigma Visibility Affects Interview Performance, Especially for the Group Expecting High Discrimination

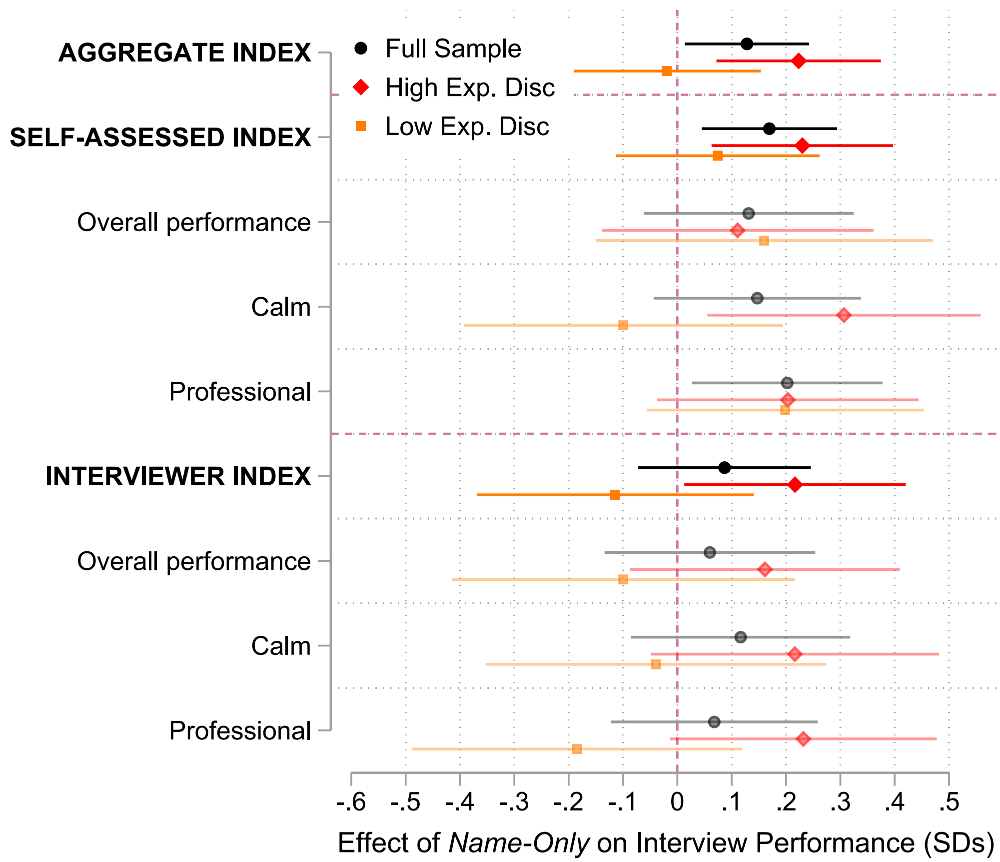

Note: The graph shows treatment effect estimates (without controls) for the full sample and conditional on expecting below or at-or-above median discrimination at baseline. The interview performance outcomes are listed on the left-hand side and described in Section 3.4. The aggregate index is an average of the self-assessed and interviewer indexes. Horizontal bars indicate 95% intervals, using robust standard errors.

Figure 7: Expected Stigma Visibility Decreases Share of Individuals in the Right Tail of Interviewer-assessed Performance Among Those Who Expect High Discrimination

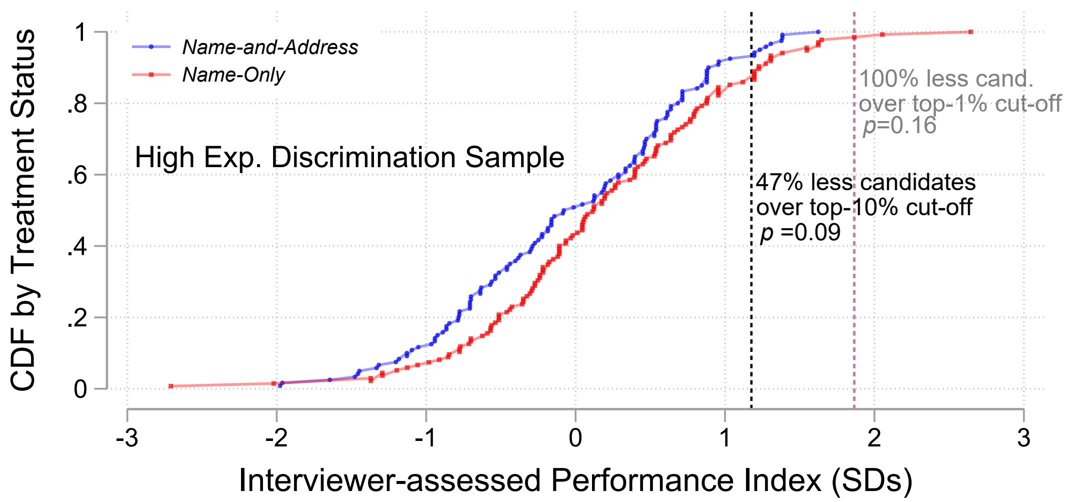

Notes: Lines show the empirical cumulative distribution of interviewer-assessed performance for each experimental condition. P-values for the decrease in representation were calculated using regression with robust SEs, without controls.

A Supporting Tables And Figures

A.1 Baseline Survey

Table A.1: Baseline Statistics

	Mean	SD	Min	Max	N
Maré resident (0/1)	0.62	0.48	0	1	2,167
Jacarezinho resident (0/1)	0.19	0.39	0	1	2,167
Manguinhos resident (0/1)	0.19	0.39	0	1	2,167
Age	25.91	6.24	18	41	2,167
Male (0/1)	0.30	0.46	0	1	2,167
White jobseeker (0/1)	0.22	0.42	0	1	2,167
Some college (0/1)	0.08	0.27	0	1	2,167
Completed regular high-school (0/1)	0.80	0.40	0	1	2,167
Working now (0/1)	0.32	0.47	0	1	2,167
Holds a formal job (0/1)	0.13	0.34	0	1	2,167
Ever worked (0/1)	0.75	0.43	0	1	2,167
Actively search last week (0/1)	0.49	0.50	0	1	2,167
Microsoft Office Experience (0/1)	0.80	0.40	0	1	1,984
Surveyor-assessed comm skills (Likert scale, 0-5)	2.79	1.10	0	4	2,167
Math test score	6.96	2.50	0	17	2,167
Heard of people refused job/fired due to address (0/1)	0.32	0.47	0	1	2,167
Believes has been refused job/fired due to address (0/1)	0.28	0.45	0	1	2,167
Own-favela expected Audit Study callback rate (%)	30.30	20.23	0	100	2,167
Adjacent non-favela expected Audit Study callback rate (%)	63.24	24.54	0	100	2,167
Racism (is reason, 0/1)	0.68	0.47	0	1	1,497
Having a different culture/speech (is reason, 0/1)	0.66	0.47	0	1	1,497
Dislike of favela residents (is reason, 0/1)	0.65	0.48	0	1	1,497
Distance to work (is reason, 0/1)	0.45	0.50	0	1	1,497
Missing days because of police raids (is reason, 0/1)	0.75	0.44	0	1	1,497
Lower skill (is reason, 0/1)	0.50	0.50	0	1	1,497
Difficulty adapting to work (is reason, 0/1)	0.47	0.50	0	1	1,497
Fear or violence (is reason, 0/1)	0.60	0.49	0	1	1,497

Note: This table presents descriptive statistics for the door-to-door baseline survey. Differences in sample sizes occur because we dropped them after introducing the Information Experiment.

Table A.2: Census (2010) Summary Statistics

Location	Population	Literate Share	White Population Share	Income per Capita in R$ (2010)
All non-favela neighborhoods in Rio	4,888,663	0.92	0.57	1376.35
All favela neighborhoods in Rio	1,391,953	0.84	0.33	382.87
Jacarezinho (favela)	37,792	0.87	0.33	349.63
Manguinhos (favela)	36,151	0.83	0.34	346.86
Maré (favela)	129,715	0.83	0.38	395.38
Bonsucesso (non-favela)	18,341	0.93	0.60	897.97
Maria da Graça (non-favela)	7,967	0.93	0.67	1126.26

Note: This table presents summary statistics from the 2010 Census.

Table A.3: Address Omission Experiment: Randomization Balance

	(1) Expects >%50 disc in audit	(2) White jobseeker (0/1)	(3) Male (0/1)	(4) Skill index	(5) Maré resident (0/1)	(6) Completed regular high schoool	(7) Working now (0/1)	(8) Ever worked (0/1)	(9) Age
Address Omission	-0.002	0.015	0.081***	0.092	-0.017	0.027	-0.005	0.053*	0.548
	(0.032)	(0.029)	(0.031)	(0.059)	(0.028)	(0.029)	(0.032)	(0.031)	(0.425)
Known Address	-0.002	-0.026	-0.004	-0.025	-0.009	0.024	-0.011	0.039	0.737*
	(0.032)	(0.028)	(0.030)	(0.058)	(0.028)	(0.029)	(0.032)	(0.031)	(0.426)
Observations	1302	1302	1302	1302	1302	1302	1302	1302	1302
Status Quo Mean	0.66	0.23	0.27	-0.07	0.80	0.76	0.33	0.69	25.35
Favela=Full *p*	1.00	0.15	0.01	0.05	0.76	0.91	0.83	0.63	0.66

* p<0.1, ** p<0.05, *** p<0.01.

Table A.4: Information Experiment: Randomization Balance

	(1) Expects >%50 disc in audit	(2) White jobseeker (0/1)	(3) Male (0/1)	(4) Skill index	(5) Maré resident (0/1)	(6) Completed regular high schoool	(7) Working now (0/1)	(8) Ever worked (0/1)	(9) Age
Favela Info	0.03	0.00	0.02	-0.02	0.04	-0.00	0.02	0.06	0.92
	(0.05)	(0.04)	(0.04)	(0.08)	(0.04)	(0.03)	(0.04)	(0.04)	(0.58)
Full Info	-0.03	0.02	0.03	-0.17**	0.04	-0.07*	-0.03	-0.01	-0.39
	(0.05)	(0.04)	(0.04)	(0.08)	(0.05)	(0.04)	(0.04)	(0.04)	(0.57)
Observations	690	690	690	690	690	690	690	690	690
No Info Mean	0.4	0.2	0.3	0.1	0.3	0.8	0.3	0.8	26.0
Favela=Full *p*	0.23	0.72	0.91	0.07	0.98	0.05	0.29	0.07	0.02

* p<0.1, ** p<0.05, *** p<0.01.

Table A.5: Interview Experiment: Randomization Balance

	(1) Expects >%50 disc in audit	(2) White jobseeker (0/1)	(3) Male (0/1)	(4) Skill index	(5) Maré resident (0/1)	(6) Completed regular high schoool	(7) Working now (0/1)	(8) Ever worked (0/1)	(9) Age
Name-Only	0.020	-0.001	0.006	0.051	-0.095**	0.019	0.060*	0.056	-0.133
	(0.048)	(0.042)	(0.043)	(0.085)	(0.047)	(0.041)	(0.033)	(0.043)	(0.565)
Observations	422	422	422	422	422	422	422	422	422
Control Mean	0.59	0.24	0.26	-0.01	0.66	0.77	0.10	0.71	25.12

* p<0.1, ** p<0.05, *** p<0.01.

Table A.6: Comparison of Samples Across the Three Experiments

Variable	(1) Address Omission Experiment N	(1) Address Omission Experiment Mean/(SE)	(2) Information Experiment N	(2) Information Experiment Mean/(SE)	(3) Interview Experiment N	(3) Interview Experiment Mean/(SE)	(1)-(2) N	(1)-(2) Mean difference	(1)-(3) Pairwise t-test N	(1)-(3) Pairwise t-test Mean difference	(2)-(3) N	(2)-(3) Mean difference
Maré resident (0/1)	1302	0.790 (0.011)	690	0.354 (0.018)	422	0.614 (0.024)	1992	0.436***	1724	0.176***	1112	-0.260***
Jacarezinho resident (0/1)	1302	0.184 (0.011)	690	0.193 (0.015)	422	0.204 (0.020)	1992	-0.009	1724	-0.020	1112	-0.011
Manguinhos resident (0/1)	1302	0.027 (0.004)	690	0.454 (0.019)	422	0.182 (0.019)	1992	-0.427***	1724	-0.156***	1112	0.271***
Age	1302	25.783 (0.174)	690	26.036 (0.236)	422	24.815 (0.283)	1992	-0.254	1724	0.967***	1112	1.221***
Male (0/1)	1302	0.295 (0.013)	690	0.303 (0.018)	422	0.265 (0.022)	1992	-0.008	1724	0.030	1112	0.037
White jobseeker (0/1)	1302	0.228 (0.012)	690	0.210 (0.016)	422	0.237 (0.021)	1992	0.018	1724	-0.009	1112	-0.027
Some college (0/1)	1302	0.064 (0.007)	690	0.080 (0.010)	422	0.071 (0.013)	1992	-0.016	1724	-0.007	1112	0.009
Completed regular high-school (0/1)	1302	0.776 (0.012)	690	0.823 (0.015)	422	0.777 (0.020)	1992	-0.047**	1724	-0.002	1112	0.046*
Working now (0/1)	1302	0.326 (0.013)	690	0.284 (0.017)	422	0.135 (0.017)	1992	0.042*	1724	0.191***	1112	0.149***
Holds a formal job (0/1)	1302	0.118 (0.009)	690	0.135 (0.013)	422	0.047 (0.010)	1992	-0.017	1724	0.071***	1112	0.087***
Ever worked (0/1)	1302	0.722 (0.012)	690	0.786 (0.016)	422	0.737 (0.021)	1992	-0.064***	1724	-0.015	1112	0.049*
Actively search last week (0/1)	1302	0.531 (0.014)	690	0.425 (0.019)	422	0.649 (0.023)	1992	0.106***	1724	-0.119***	1112	-0.225***
Surveyor-assessed comm skills (Likert scale, 0-5)	1302	2.795 (0.029)	690	2.797 (0.045)	422	3.001 (0.051)	1992	-0.002	1724	-0.206***	1112	-0.204***
Math test score	1302	6.960 (0.072)	690	6.945 (0.091)	422	6.919 (0.115)	1992	0.015	1724	0.041	1112	0.026
Reservation wage (USD)	1301	253.155 (3.016)	690	246.173 (3.215)	422	231.962 (2.736)	1991	6.983	1723	21.193***	1112	14.211***

* p<0.1, ** p<0.05, *** p<0.01.

Table A.7: Effects of Information on Beliefs for Under- and Overestimators

	(1) Disc (%)	(2) Cb. Own Neigh (%)	(3) Cb. Own Neigh (%)	(4) Cb. Own Neigh (%)	(5) Cb. Other Neigh (%)	(6) Cb. Other Neigh (%)	(7) Cb. Other Neigh (%)
Favela Info	-7.288** (3.487)	-2.419 (1.886)	-3.344* (1.961)	2.995 (3.925)	-8.049*** (2.149)	-8.521*** (2.116)	3.337 (9.814)
Full Info	-20.371*** (3.692)	-1.428 (1.973)	-3.962* (2.055)	11.417** (4.541)	-13.923*** (2.217)	-14.918*** (2.193)	12.625 (12.195)
Observations	690	690	554	136	690	637	53
Sample	All	All	Overestimators	Underestimators	All	Overestimators	Underestimators
Control Mean	35.46	36.60	40.62	18.36	59.29	61.34	29.71
Favela=Full p	0.00	0.60	0.76	0.05	0.01	0.00	0.40

* p<0.1, ** p<0.05, *** p<0.01.

Table A.8: Effects on Interview Performance for the High Expected Discrimination and White Individuals Are Robust to Including Other Interacted Covariates

	(1) Interviewer-assessed performance	(2) Interviewer-assessed performance	(3) Interviewer-assessed performance	(4) Interviewer-assessed performance	(5) Self-assessed performance	(6) Self-assessed performance	(7) Self-assessed performance	(8) Self-assessed performance
Name-Only	-0.11	0.02	-0.18	0.44	0.07	0.12*	0.03	0.64*
	(0.13)	(0.09)	(0.14)	(0.46)	(0.09)	(0.07)	(0.10)	(0.37)
Name-Only × High Exp. Disc.	0.33**		0.32*	0.24	0.16		0.16	0.14
	(0.17)		(0.16)	(0.15)	(0.13)		(0.13)	(0.13)
Name-Only × White jobsee.		0.29	0.29	0.31*		0.19	0.19	0.17
		(0.20)	(0.20)	(0.19)		(0.15)	(0.15)	(0.15)
Observations	422	422	422	422	422	422	422	422
Other Interactions?	No	No	No	Yes	No	No	No	Yes

* p<0.1, ** p<0.05, *** p<0.01. We estimate $y_i = \alpha + \beta \text{Name-Only}_i + \sum_j \gamma^j \text{Name-Only} \times X_i^j + \nu^j X_i^j$ for different sets of covariates. Other Interactions include skill, having completed regular high school, employment status. having ever worked, age, and a dummy for having reservation wage over R$1,500.

Table A.9: Table 2 with Double-Lasso Controls

	(1) Agg. Performance Index (SDs)	(2) Applied (%)	(3) Showed Up (%)	(4) Applied (%)	(5) Showed Up (%)
$\downarrow \mathbb{E}$[Disc.] × white	0.29***	7.47	8.74*	-11.89	-17.96***
	(0.11)	(5.81)	(5.08)	(8.29)	(6.20)
$\downarrow \mathbb{E}$[Disc.] × non-white	0.05	-7.37**	-1.17	6.07	1.94
	(0.06)	(3.20)	(2.51)	(4.24)	(3.46)
Observations	422	1302	1302	690	690
Treatment	$\downarrow \mathbb{E}$[address visibility]	$\downarrow \mathbb{E}$[address visibility]	$\downarrow \mathbb{E}$[address visibility]	$\downarrow \mathbb{E}$[mkt-level disc.]	$\downarrow \mathbb{E}$[mkt-level disc.]
Control Average	-0.0	44.1	18.2	42.3	20.7
White=non-white p-value	0.05	0.03	0.08	0.05	0.01

Table A.10: Two-stage Least Squares Estimates of The Effect of the Expected Callback Rate and Anti-Favela Discrimination on Application Decisions

	(1) Applied (%)	(2) Show Up (%)	(3) Obfuscated in application (%)
Posterior Expected Callback for Own Favela (%)	-0.45	-0.14	-0.55*
	(0.51)	(0.41)	(0.33)
Posterior Expected Discrimination Rate (%)	-0.11	0.05	0.11
	(0.21)	(0.18)	(0.14)
Observations	690	690	690
No Info Mean	39.8	19.9	10.6

Note: This table uses variation in beliefs induced by the information treatments *Favela Info* and *Full Info* to estimate their effects on application decisions. Instrumented variables are the expected callback rate the HR firm would implement in the person's favela of residence and the implied discrimination rate (percent drop in callback caused by being from the favela instead of living just outside it). Instruments are the the treatment assignment interacted with i) dummy for overestimating the favela callback rate when predicting the audit study, ii) prediction error when predicting that callback rate, iii) dummy for overestimating the discrimination in callbacks when predicting the audit study iv) prediction error in predicting that discrimination rate. Outcomes are completing the online application form, attending the interview, and obfuscating address in the online application form.

Figure A.1: Figure 4 with Lasso-selected Controls

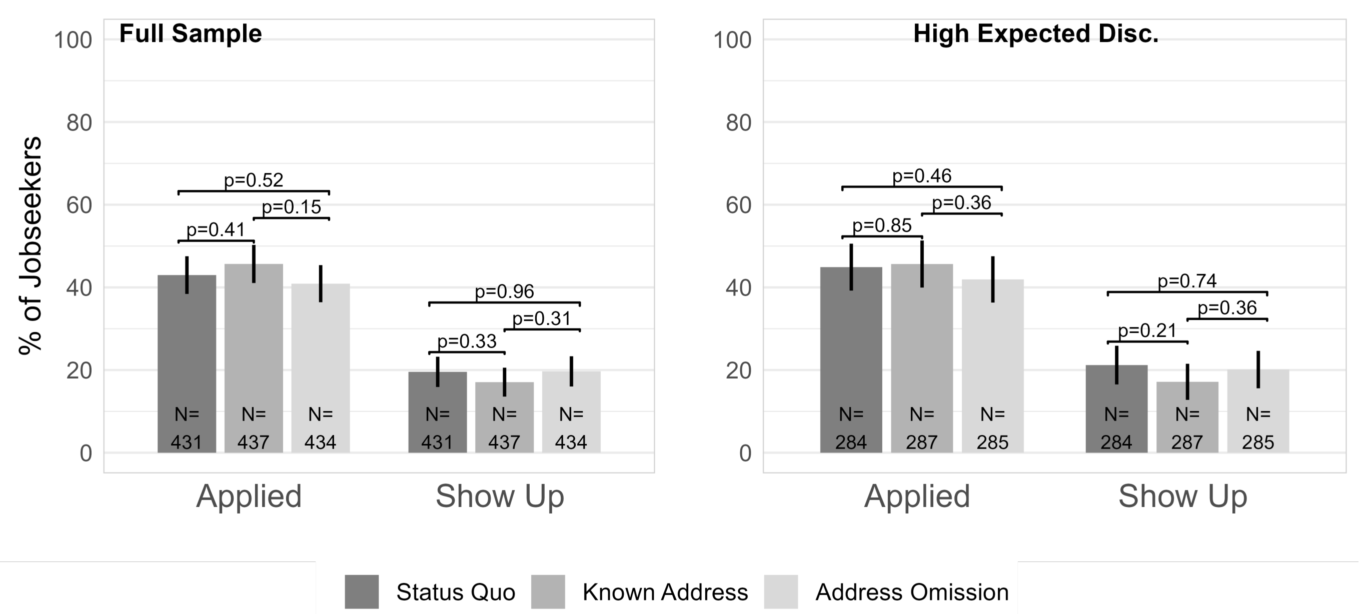

Figure A.2: Figure 5 with Lasso-selected Controls

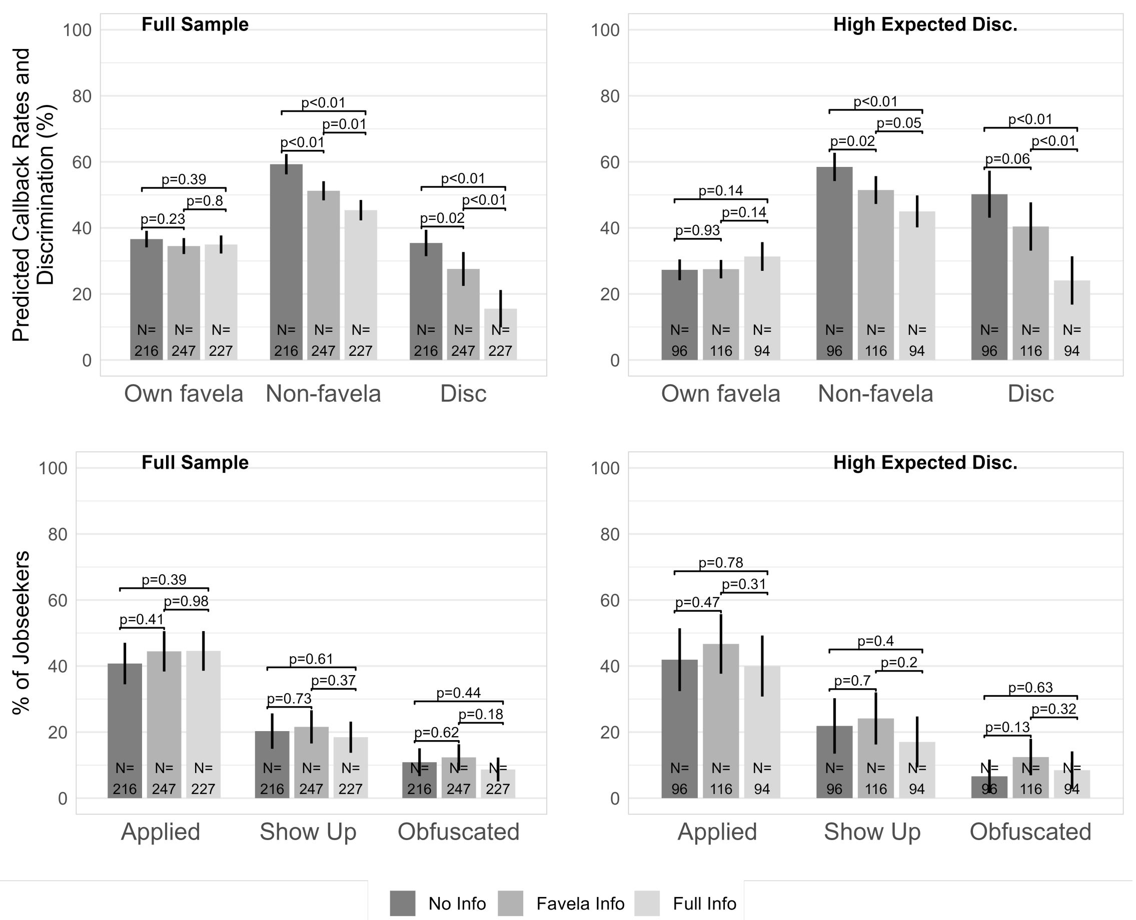

Figure A.3: Figure 6 with Lasso-Selected Controls

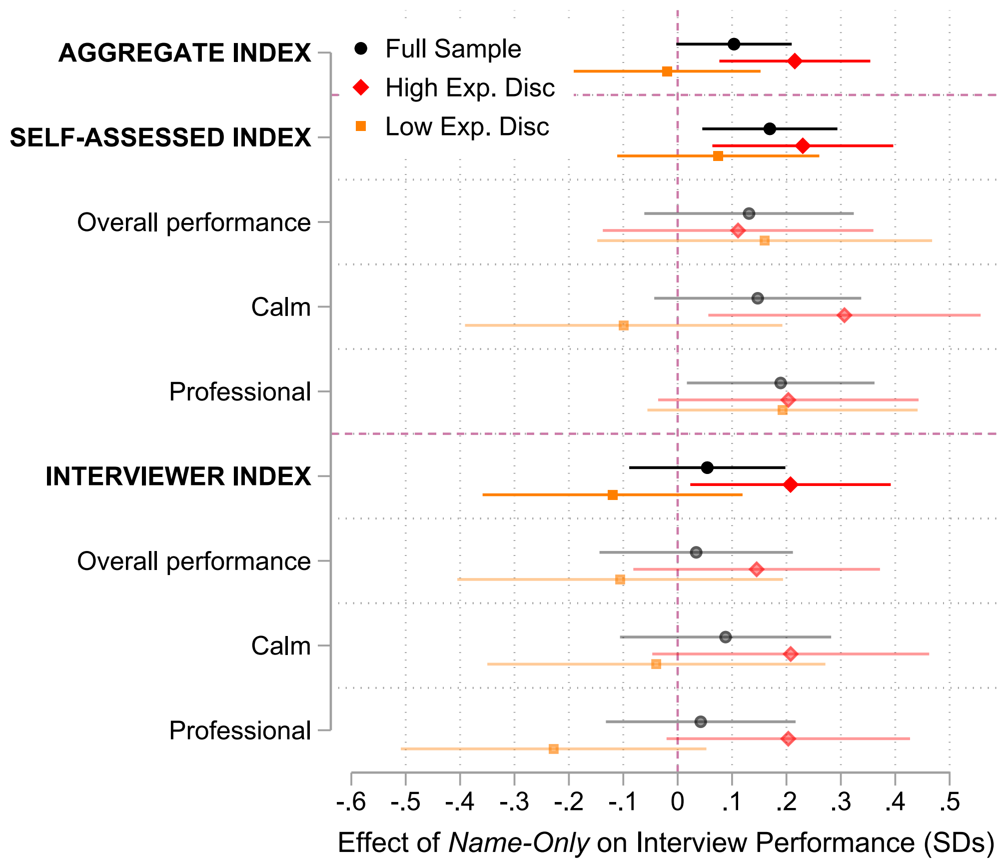

Note: The graph shows treatment effect estimates (using double-lasso selected controls) for the full sample and conditional on expecting below or at-or-above median discrimination at baseline. The interview performance outcomes are listed on the left-hand side and described in Section 3.4. The aggregate index is an average of the self-assessed and interviewer indexes. Horizontal bars indicate 95% intervals, using robust standard errors.

Figure A.4: Effects of Information Treatments on Beliefs and Applications by Whether Job-seekers Initially Under- or Overestimated the Favela Callback Rate

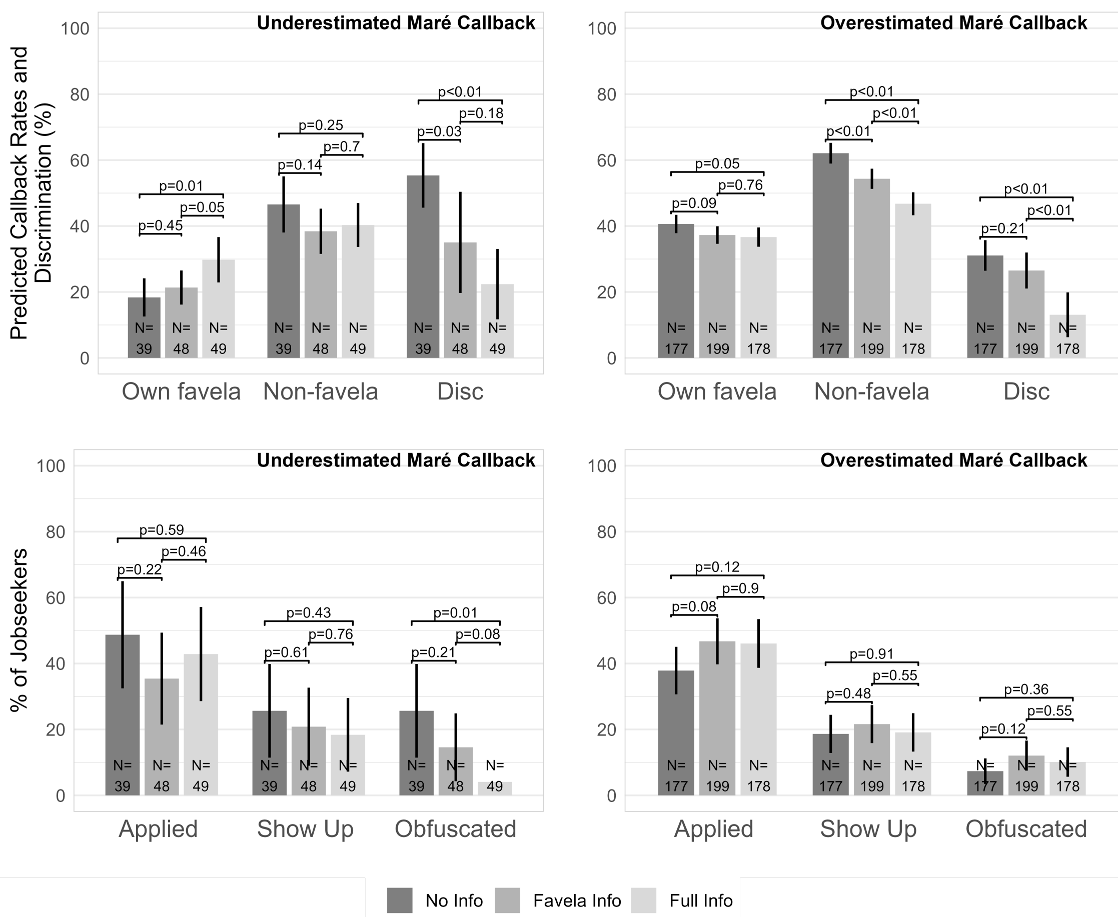

Figure A.5: Heterogeneous Effects in the Address Omission Experiment – No Controls

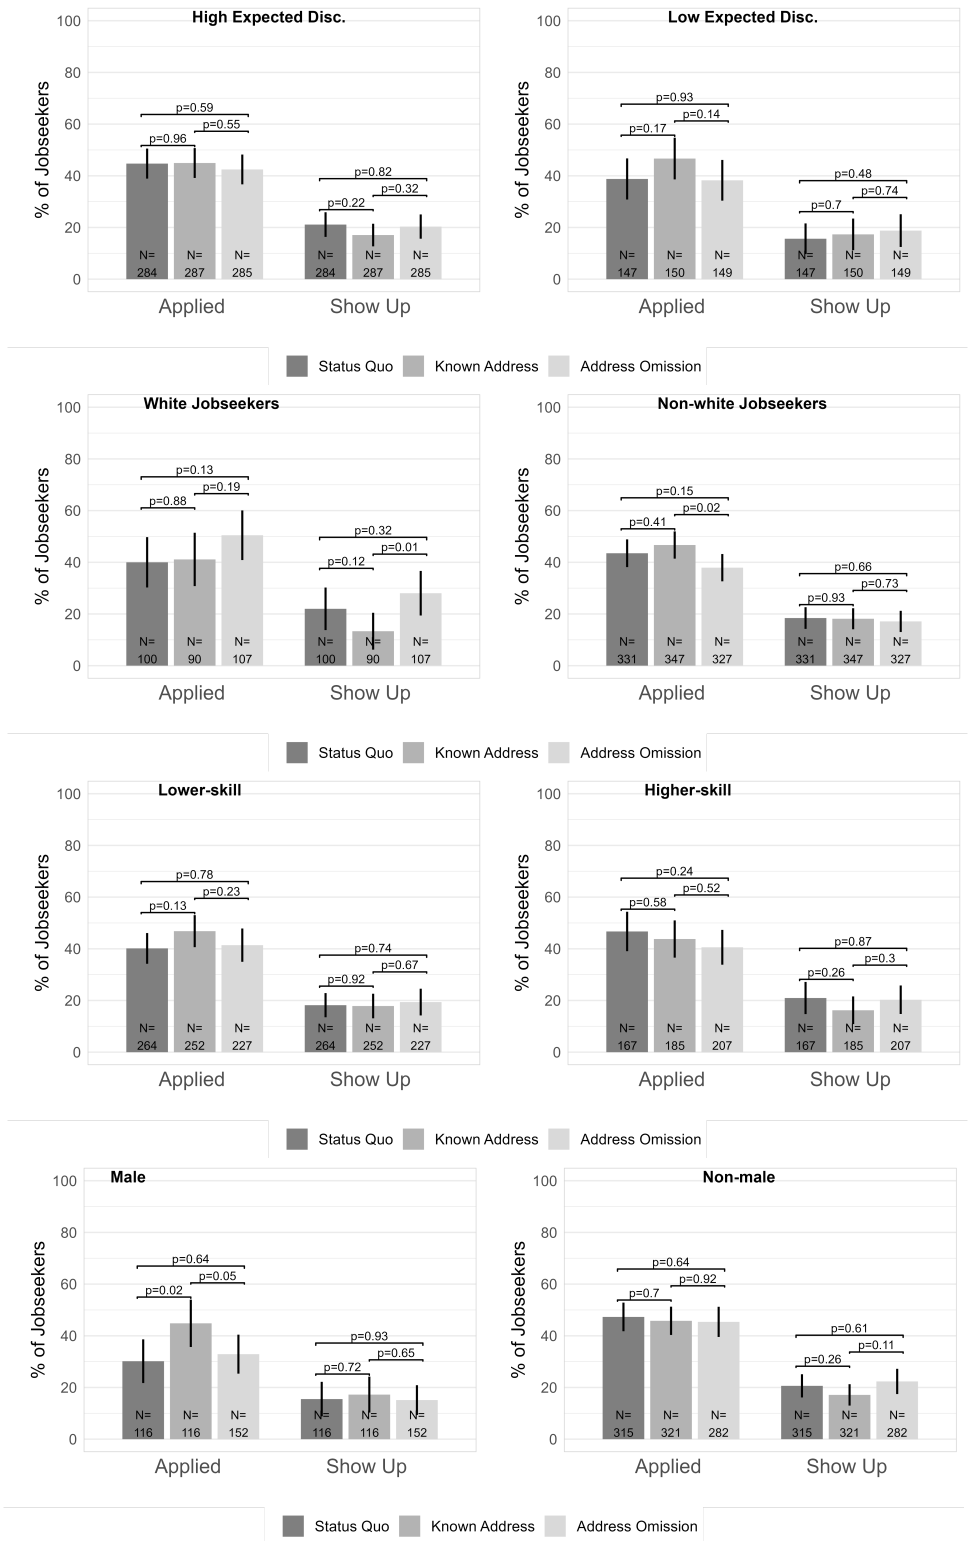

Figure A.6: Heterogeneous Effects in the Address Omission Experiment – Double-lasso Controls

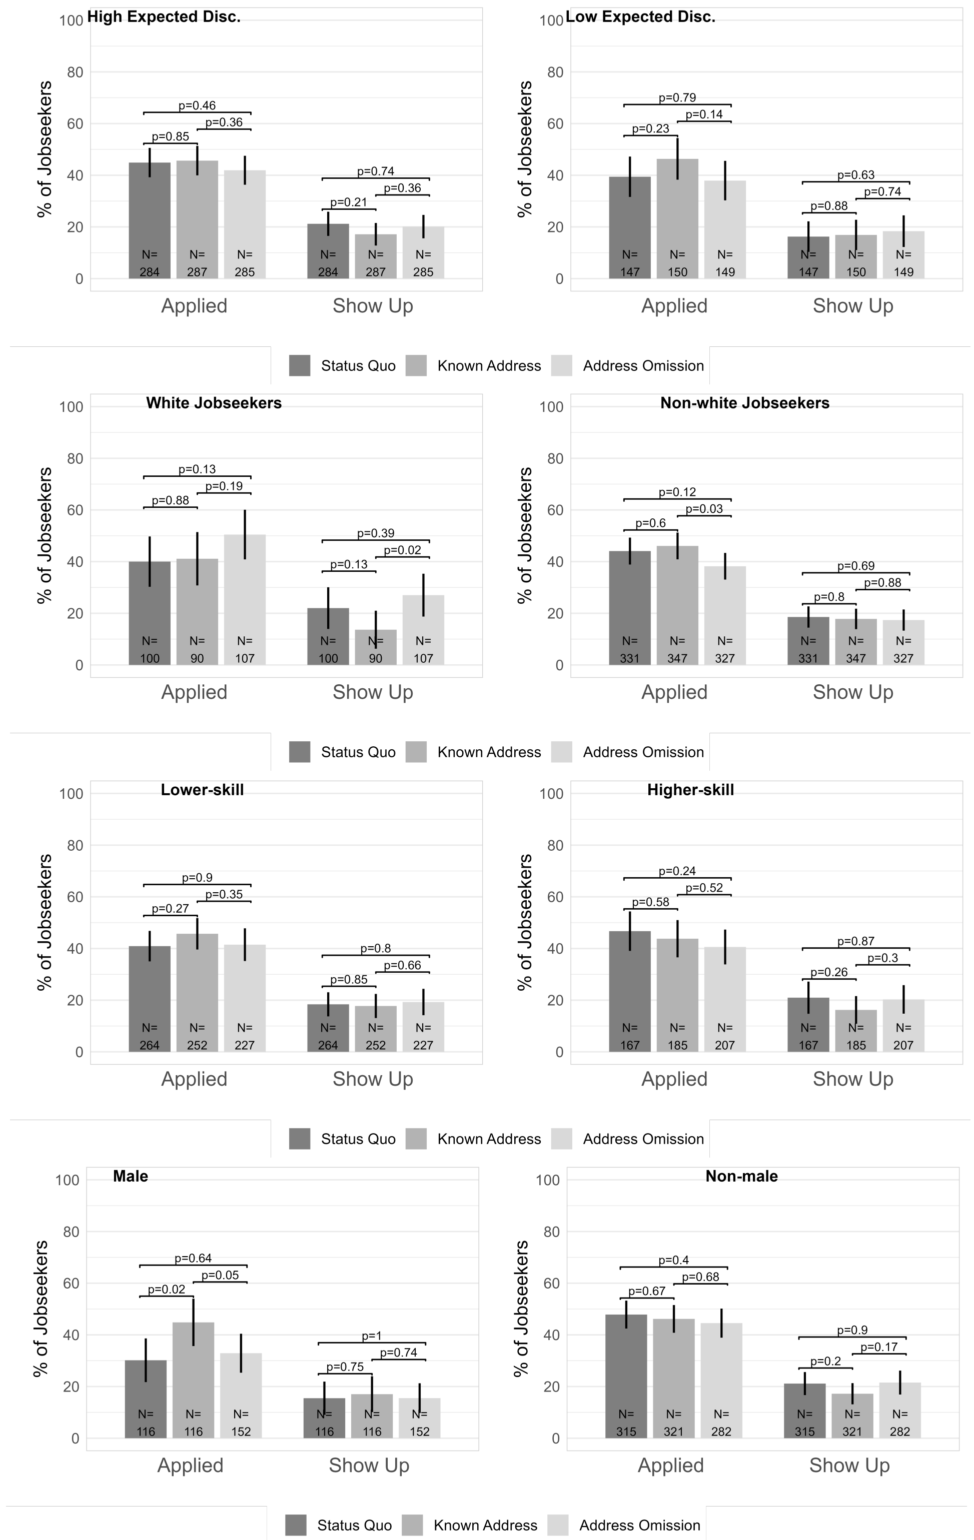

Figure A.7: Heterogeneous Effects of Information Treatments – No Controls

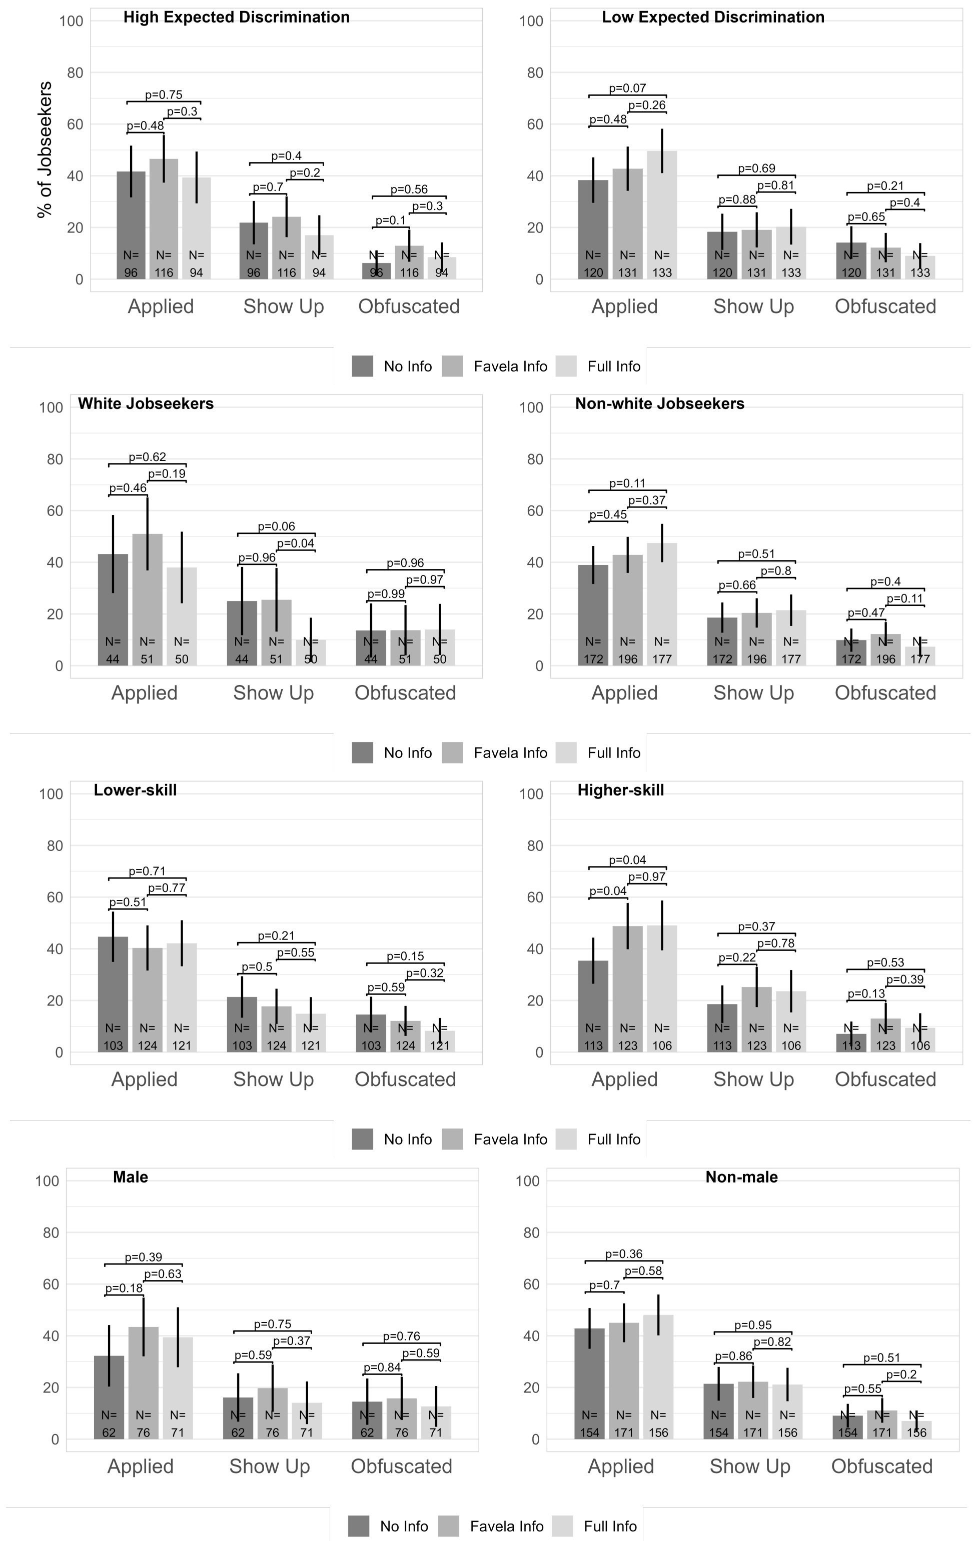

Figure A.8: Heterogeneous Effects of Information Treatments – Double-lasso Controls

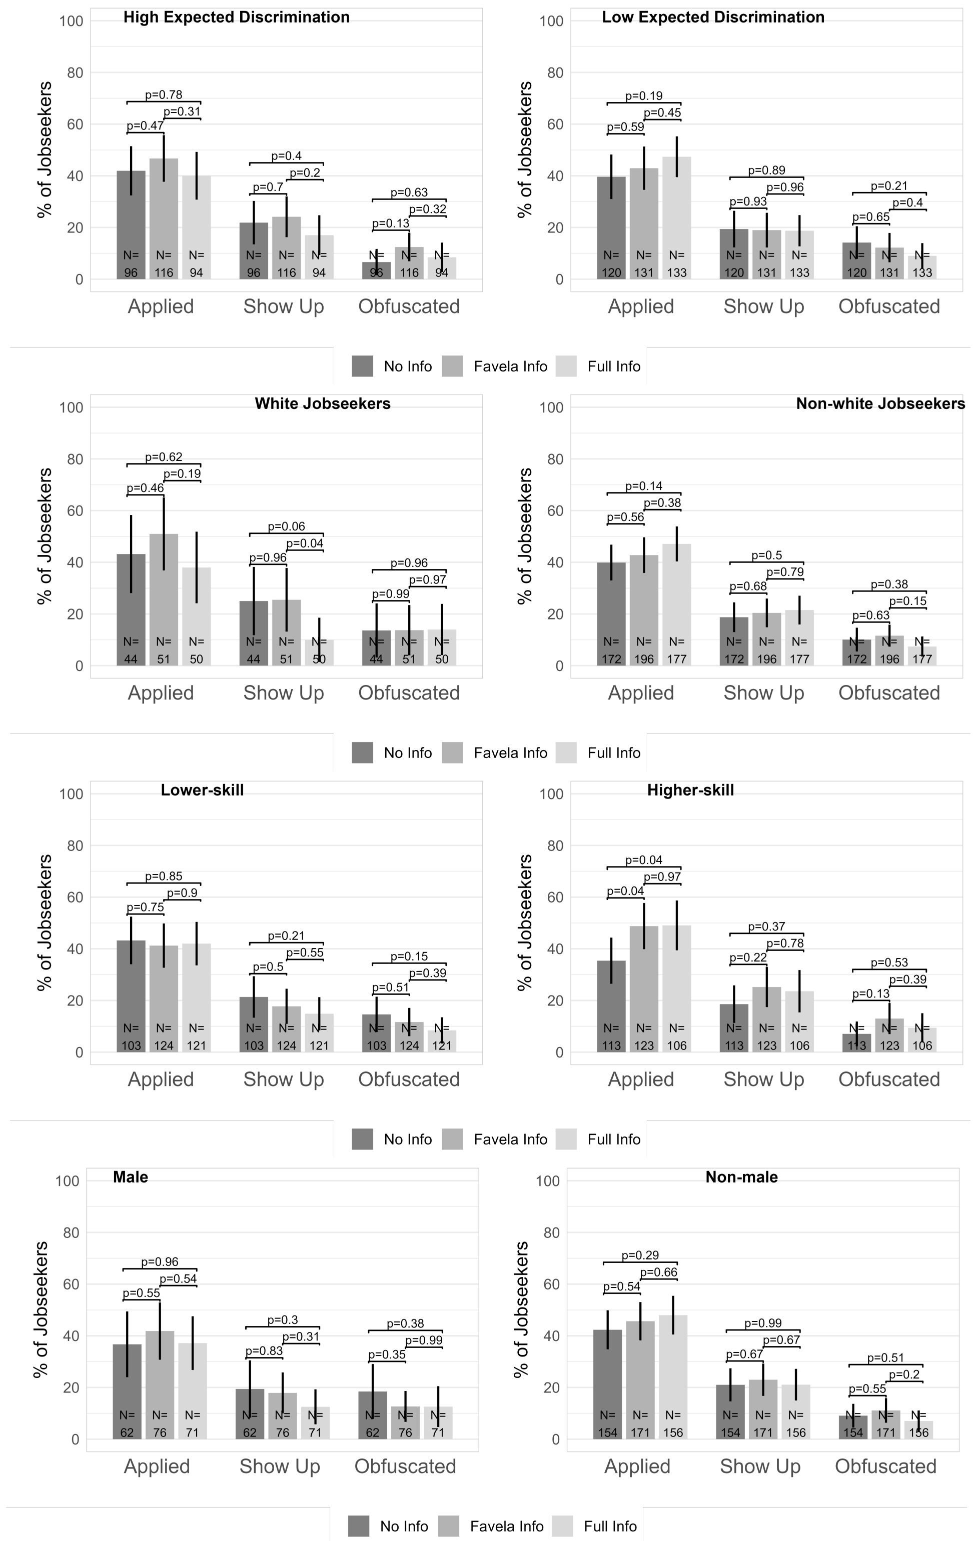

Figure A.9: Heterogeneous Effects of *Name-Only*

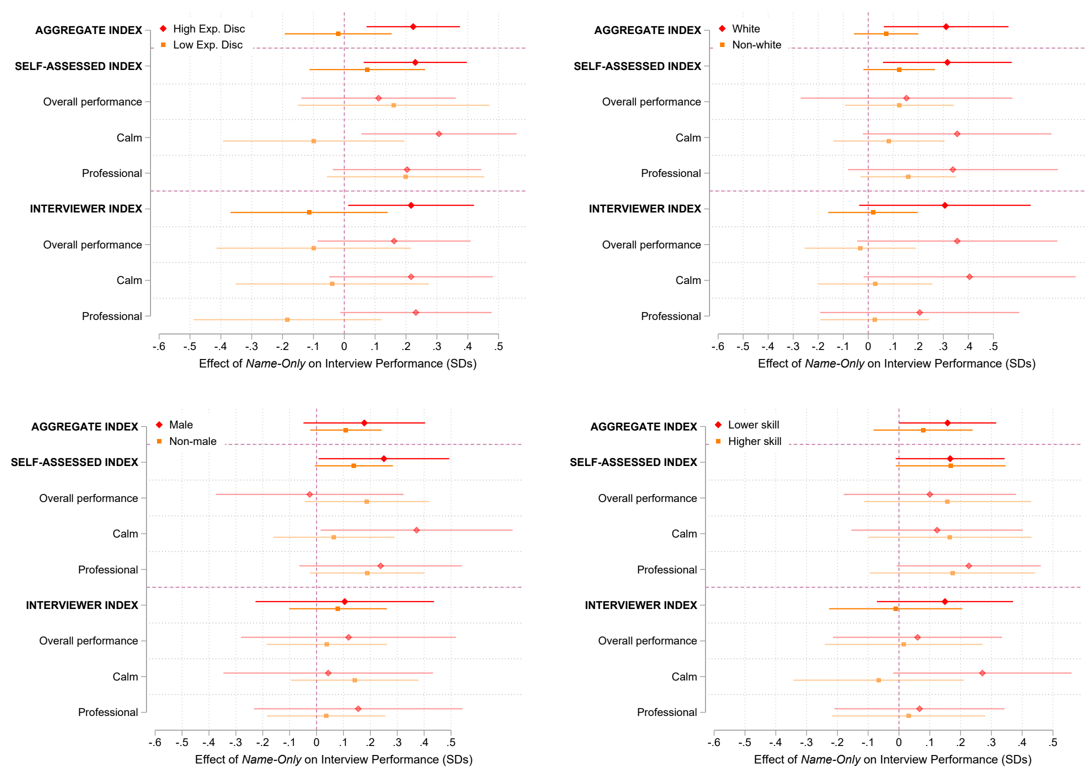

Figure A.10: Predicted *vs.* Actual Discrimination Rates

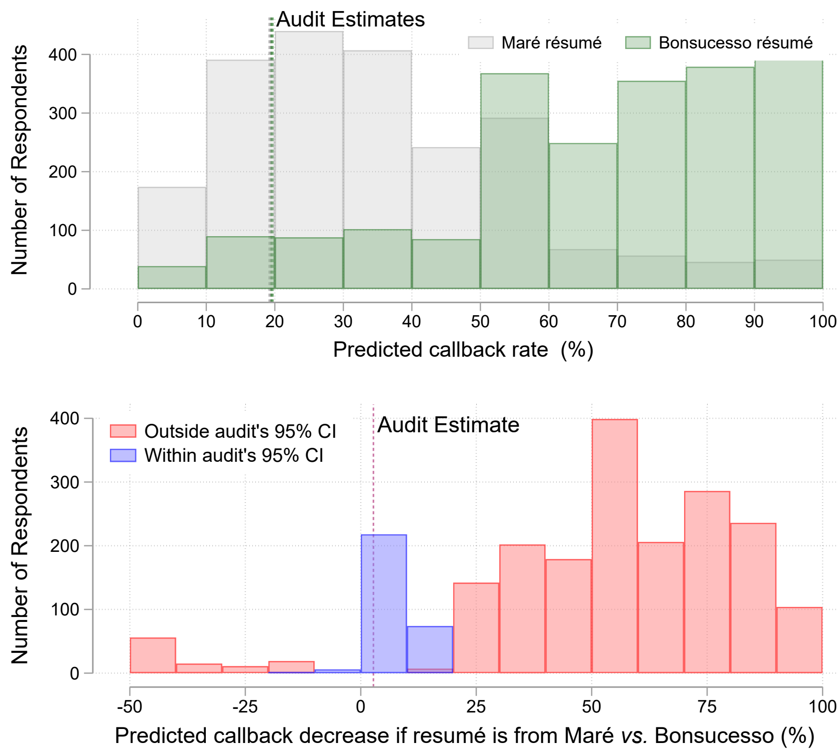

Figure A.11: Predicted Audit Study Discrimination Correlates with Other Measures of Expected Discrimination

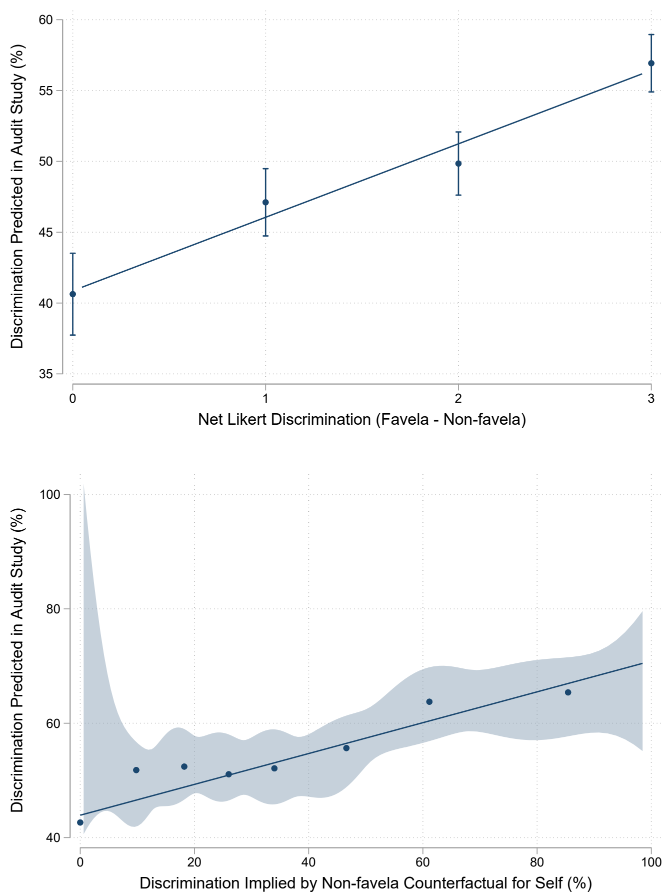

Note: Negative values of discrimination are pooled with zero discrimination – since there are few observation with negative discrimination, which make estimates noisy. We construct the Likert discrimination measure by taking the Liker-scale answers of how much employers discriminate against individuals in each neighborhood (from no discrimination to a lot), converting them into ordered integers, and and taking the difference. We calculate the discrimination for the counterfactual self by comparing the beliefs about one's job-finding probability over the next six months to "somebody just like you, but from [adjacent non-favela]".

Figure A.12: Belief Update in Information Experiment Occurs for Maré and Non-Maré Residents

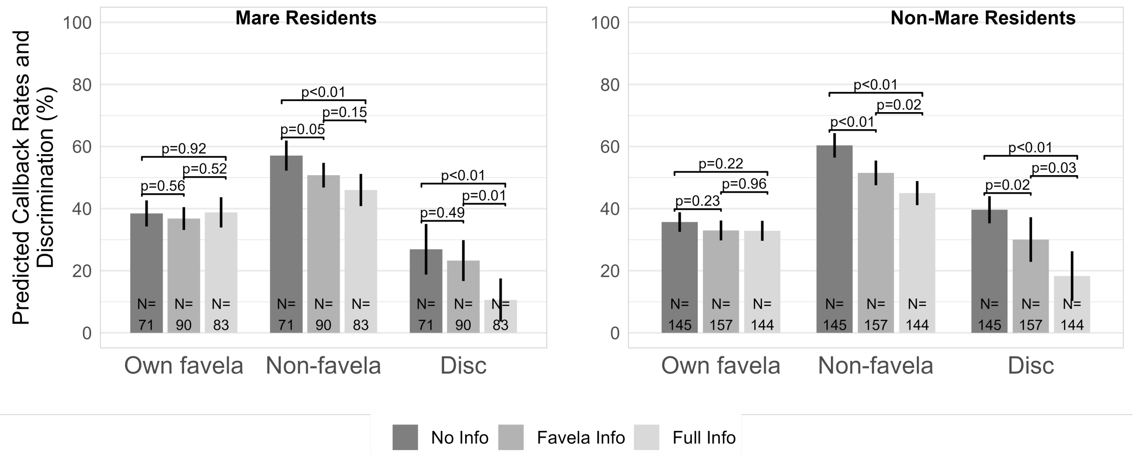

Figure A.13: Cumulative Distribution of the Interviewer-assessed Performance by Experimental Condition – Whole Sample

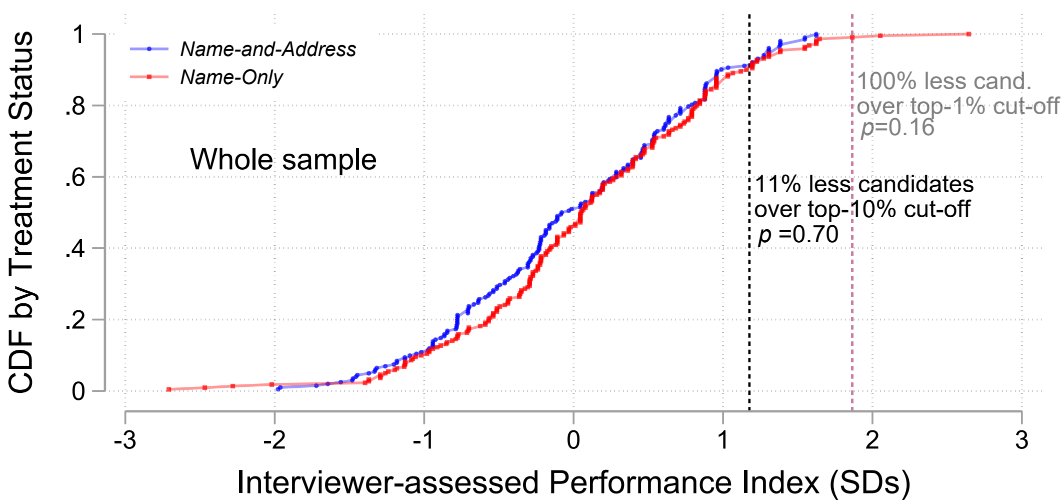

Notes: Lines show the empirical cumulative distribution of interviewer-assessed performance for each experimental condition. P-values for the decrease in representation were calculated using regression with robust SEs, without controls.

Figure A.14: Cumulative Distribution of the Interviewer-assessed Performance by Experimental Condition – White Jobseekers

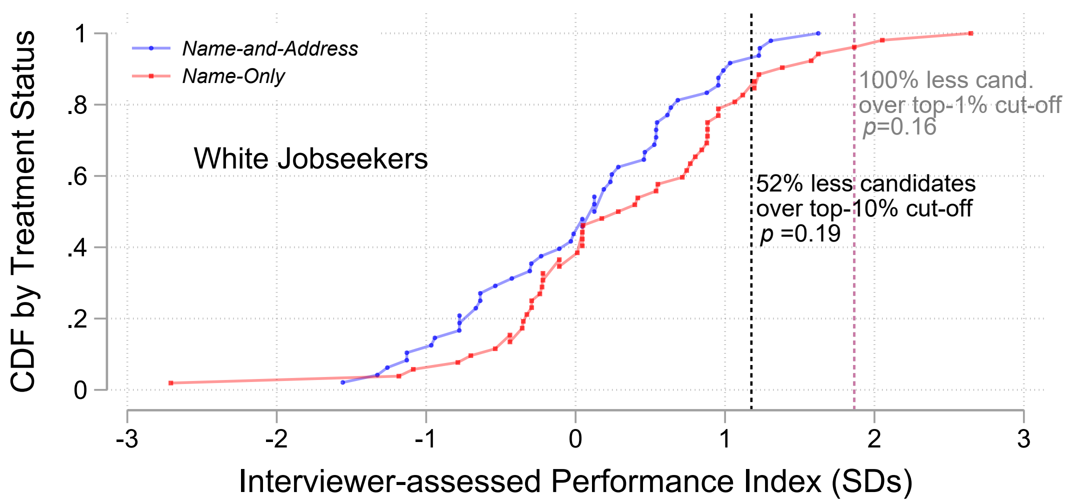

Notes: Lines show the empirical cumulative distribution of interviewer-assessed performance for each experimental condition. P-values for the decrease in representation were calculated using regression with robust SEs, without controls.

B Deviations from the Pre-Analysis Plan

- We initially planned to stratify the randomization in the Interview Experiment by predicted discrimination level *and* previous treatment assignments. On implementation we kept only stratification by the discrimination level. That is because, given the logistical constraints and lower-than-expected interview show-up rates, the batch sizes for the interview stage would generate a very small number of observation per strata.
- We pre-registered our in-survey math test as the main skill measure, but we later judged it was too narrow with respect to a sales job. Hence, we also included education and a measure of communication skills.
- The receptionist randomized the treatment of ten participants at the office, and results do not change by excluding them. She conducted the on-the-spot randomization when either i) she could not locate the jobseeker's treatment status (e.g., due to internet connection issues), or ii) a candidate was mistakenly invited to the interview before being assigned a randomization batch, or iii) the number of candidates schedules for a period was too low for make up a single strata.
- CHOICE OF THE MAIN EXPECTED DISCRIMINATION MEASURE – WE DROPPED SOME OF THE OTHER MEASURES WHEN INTRODUCING THE AUDIT STUDY

C Audit Study

Picking Résumé Addresses. For addresses in each neighborhood, we picked streets that were i) entirely contained in the neighborhood, ii) in the postal office list, and iii) up to a 15-minute walk from a bus stop in the avenue between Maré and Bonsucesso. These choices guaranteed that employers could back out neighborhood unambiguously, and keep commuting time to any job as constant as possible.

Table C.1: Audit Study Results

	(1) Callback (%)	(2) Callback (%)	(3) Callback (%)
Maré résumé	-0.34	-0.40	-1.04
	(1.28)	(1.29)	(1.18)
Observations	1400	1400	1400
No Info Mean	16.96	16.96	16.96
Controls	No	Yes	No
Job FEs	No	No	Yes

Note: Outcome variable evaluates to 100 if the application received a positive response and zero otherwise. Maré résumé is a dummy for the fictitious applicant being from Maré. Controls include the job's city region, and the website in which we found it. The callback level here is about 3% lower than than the numbers used in the Information Experiment because for the regressions we only consider callbacks we could link to unique postings. Standard errors clustered at the posting level shown between parenthesis.

C.1 Audit Study Neighborhoods

Figure C.1: Bonsucesso (Non-Favela) *vs.* Maré (Favela)

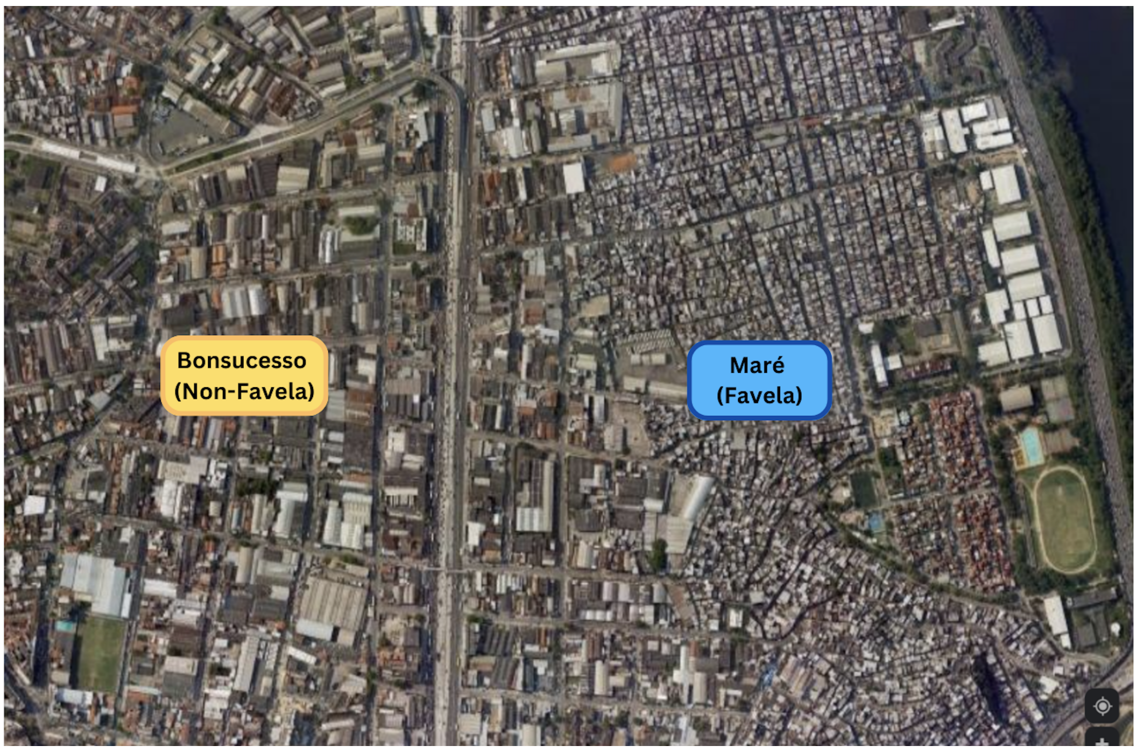

Note: This image shows the geographic location of the two neighborhoods for the audit study: Bonsucesso (Non-Favela) and Maré (Favela). The large avenue in the picture is the divide between each region.

Figure C.2: Example Résumé – Maré home address

ROBSON DE FREITAS

30 YEARS OLD • BRAZILIAN • SINGLE

CONTACT

(21) 99878-2186

guilhermeantonioalmeida3@gmail.com

Carlos Lacerda Street, 102 - Maré, RJ

EDUCATION

CE Olga Benário Prestes
High School. Full time.
feb. 2008 - dec. 2010

SENAC
Logistics Tecnitian.
feb. 2011 - dec. 2011

COMPLEMENTARY COURSES

Customer Service
SEBRAE - 2012

Customer Success
SEBRAE - 2014

Sales Management
FGV - 2016

LANGUAGES

Intermediate english.

SKILLS

Clear and objective communication; Proactivity;
Empathy;
Focus on results.

ADDITIONAL INFORMATION

Available for work on weekends.

OBJECTIVE

Sales professional with 12 years of experience. I seek new opportunities for professional growth in a collaborative work environment.

WORK EXPERIENCE

Hering

Salesperson (sep. 2021 - oct. 2022)

- Direct customer service
- Guide the customer on product specifications

Aviator

Salesperson (aug. 2016 - jun. 2021)

- Responsible for controlling the demand of orders in stock
- Assist customers in choosing products

Loja Del Rey

Salesperson (nov. 2015 - may 2016)

- Direct customer service
- Guide the customer on product specifications

Di Santinni

Sales assistant (jul. 2014 - jun. 2015)

- Responsible for controlling the demand of orders in stock
- Assist customers in choosing products

Cashier (aug. 2013 - jun. 2014)

- Act directly in customer service, finalizing the purchase and issuing the invoice

General Autopeças

Shop assistant (oct. 2011 - mar. 2013)

- Guiding customers in choosing and demonstrating how products work

Loja Impecável

Cashier (jan. 2011 - aug. de 2011)

- Opening and closing the cash register
- Responsible for processing and receiving payment

Note: This image shows one of the résumés used in the audit study. We drew the red box around the address in this picture for emphasis. It was not present in the original résumé.

Figure C.3: Example Résumé – Bonsucesso Address

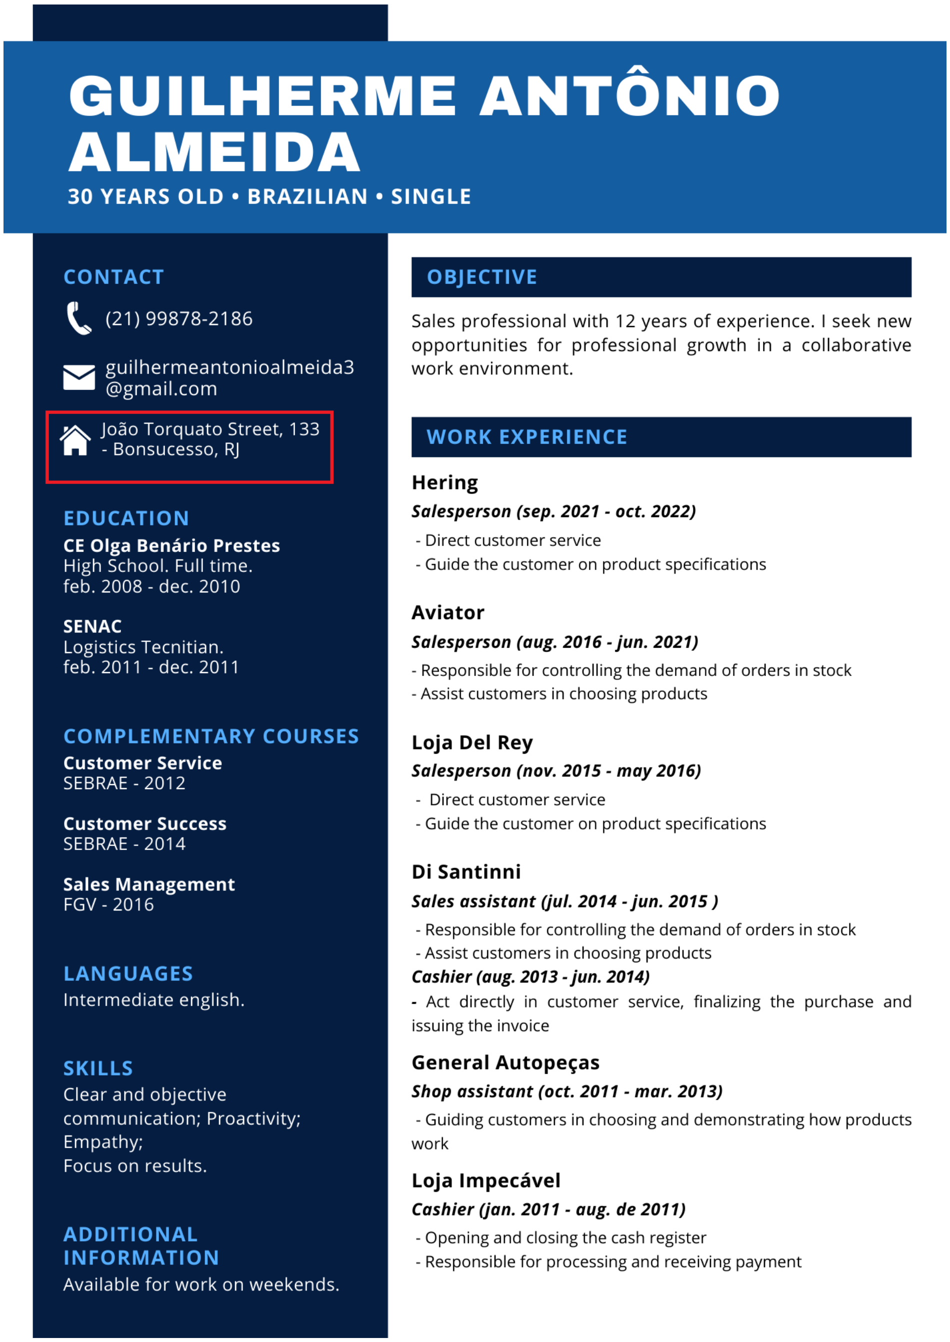

GUILHERME ANTÔNIO ALMEIDA

30 YEARS OLD • BRAZILIAN • SINGLE

CONTACT

(21) 99878-2186

guilhermeantonioalmeida3
@gmail.com

João Torquato Street, 133
- Bonsucesso, RJ

EDUCATION

CE Olga Benário Prestes
High School. Full time.
feb. 2008 - dec. 2010

SENAC
Logistics Tecnitian.
feb. 2011 - dec. 2011

COMPLEMENTARY COURSES

Customer Service
SEBRAE - 2012

Customer Success
SEBRAE - 2014

Sales Management
FGV - 2016

LANGUAGES

Intermediate english.

SKILLS

Clear and objective communication; Proactivity;
Empathy;
Focus on results.

ADDITIONAL INFORMATION

Available for work on weekends.

OBJECTIVE

Sales professional with 12 years of experience. I seek new opportunities for professional growth in a collaborative work environment.

WORK EXPERIENCE

Hering
Salesperson (sep. 2021 - oct. 2022)
- Direct customer service
- Guide the customer on product specifications

Aviator
Salesperson (aug. 2016 - jun. 2021)
- Responsible for controlling the demand of orders in stock
- Assist customers in choosing products

Loja Del Rey
Salesperson (nov. 2015 - may 2016)
- Direct customer service
- Guide the customer on product specifications

Di Santinni
Sales assistant (jul. 2014 - jun. 2015)
- Responsible for controlling the demand of orders in stock
- Assist customers in choosing products

Cashier (aug. 2013 - jun. 2014)
- Act directly in customer service, finalizing the purchase and issuing the invoice

General Autopeças
Shop assistant (oct. 2011 - mar. 2013)
- Guiding customers in choosing and demonstrating how products work

Loja Impecável
Cashier (jan. 2011 - aug. de 2011)
- Opening and closing the cash register
- Responsible for processing and receiving payment

Note: Image shows one of the résumés used in the audit study. We drew the red box around the address in this picture for emphasis. It was not present in the original résumé.

D Supporting Materials

Figure D.1: Door-to-Door Baseline Survey
Notes: This Figure shows surveyors interviewing research participants in Maré.

Figure D.2: Predicted Discrimination Baseline Script

Now I'm going to ask you some questions about the differences between job seekers from different neighborhoods. We know the right answer to two of them.

If, at the end of our project, you are among the 10 people who came closest to getting these two questions right, you will receive an additional $100 Brazilian reais.

Let me tell you the story to start:

At the beginning of our project, the researchers organizing this study heard from the population of several favelas here in Rio about how it was more difficult to apply for a formal job living in a community.

To really understand the size of the challenge, researchers sent 1,400 applications with fake résumes, but as if they were real people, for 700 vacancies in sales in the city of Rio.

The résumes were from men and women, from people with experience and suitable for each vacancy.

The only difference between the résumes was that some said that the address was from Bonsucesso, and others said that the address was from Maré.

I will give you a moment to look at an example of one of the resumes sent.

[PASS THE TABLET TO THE INTERVIEWEE]

The researchers calculated WHAT PERCENTAGE of résumes sent with BONSUCESSO's address were selected (for example, for a training period) or called for an interview.

They also calculated this percentage for MARE's résumes.

To get the additional $100 Brazilian reais, I'm going to ask you to try to guess what they found, okay?

- *WHAT PERCENTAGE of résumes sent with BONSUCESSO's address do you guess were selected or invited for an interview?*
- AND WHAT PERCENTAGE OF MARE's?

Note: This Figure displays how we elicited prior beliefs about discrimination against favela dwellers.

Figure D.3: Job Descriptions

Job Description 1
Sales Representative

Duties:
- Trial, sale, and demo of products, focusing on customer satisfaction and loyalty, ensuring the cleanliness and organization of the store

Prerequisites:
- High School Degree
- Skills with persuasion and approaching
- Office Package

Desirable:
- Have a good beauty repertoire (knowing products, competing brands and influencers);
- Results-oriented

Job Description 2
Direct Sales Promoter

Duties:
- Responsible for attracting and prospecting new resellers in its operating unit. Fill out forms, register and deliver the documentation to the Direct Sale center.

Prerequisites:
- High School Degree

Desirable:
- Experience with negotiation and persuasion to charm customers

Job Description 3
Direct Sales Supervisor

Duties:
- Responsible for receiving payments, operating sales systems, issuing invoices, making contact with resellers

Prerequisites:
- High School Degree
- Office Package
- Experience with sales and payments

Desirable:
- Ability to do math
- Good verbal and written communication
- Detail oriented

Note Job descriptions as presented in the online application forms.

Figure D.4: Second Screen of the Application Form of Each Experimental Condition in the Address Omission Experiment

To apply for these jobs, you must complete this form. The UBC research team has sent us only your NAME and MOBILE PHONE.

Please confirm your NAME:

- My name is {e://Field/name}
- No, my correct NAME is:

Please confirm your MOBILE PHONE/WHATSAPP:

- My MOBILE PHONE/WHATSAPP is the one in which I received the link to this form
- No, my correct MOBILE PHONE/WHATSAPP is:

We also need your home address:

Street

Number and unit if applicable

Neighborhood -- start typing and select your neighborhood

(a) *Status Quo*

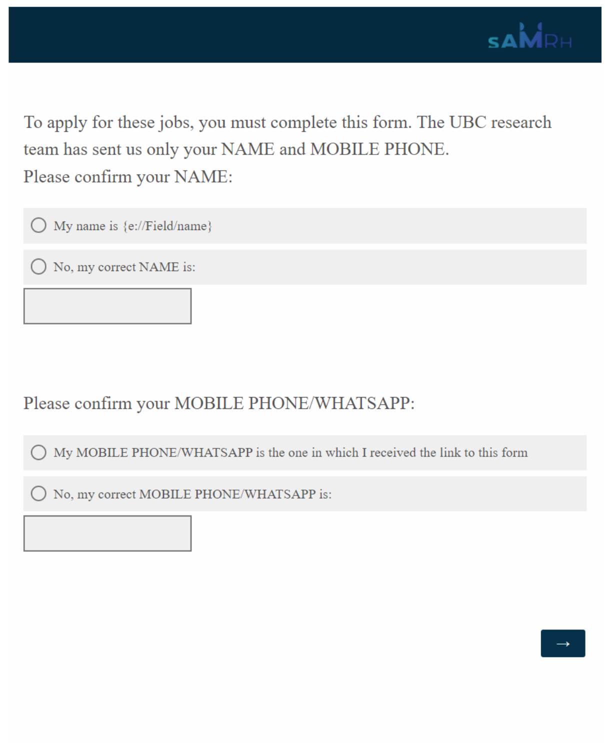

To apply for these jobs, you must complete this form. The UBC research team has sent us only your NAME and MOBILE PHONE.
Please confirm your NAME:

- My name is {e://Field/name}
- No, my correct NAME is:

Please confirm your MOBILE PHONE/WHATSAPP:

- My MOBILE PHONE/WHATSAPP is the one in which I received the link to this form
- No, my correct MOBILE PHONE/WHATSAPP is:

→

(b) *Address Omission*

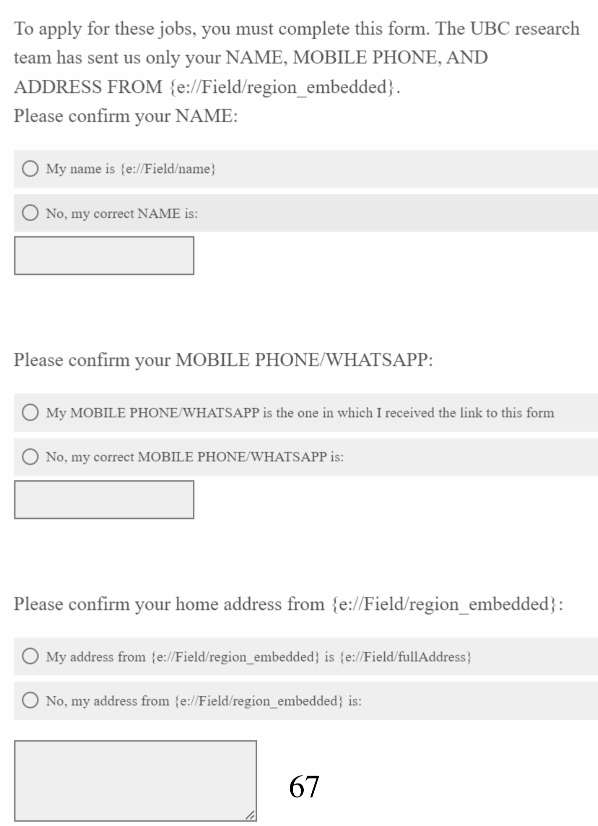

To apply for these jobs, you must complete this form. The UBC research team has sent us only your NAME, MOBILE PHONE, AND ADDRESS FROM {e://Field/region_embedded}.
Please confirm your NAME:

- My name is {e://Field/name}
- No, my correct NAME is:

Please confirm your MOBILE PHONE/WHATSAPP:

- My MOBILE PHONE/WHATSAPP is the one in which I received the link to this form
- No, my correct MOBILE PHONE/WHATSAPP is:

Please confirm your home address from {e://Field/region_embedded}:

- My address from {e://Field/region_embedded} is {e://Field/fullAddress}
- No, my address from {e://Field/region_embedded} is:

(c) *Known Address*

(a) Co-Working Reception (b) Interview Room

Figure D.5: Interview Co-Working Space

D.1 Interview Script

You [the interviewer] must treat all candidates equally and as uniformly as possible. Ideally, your tone will be friendly and reserved.

Introduce yourself and confirm the candidate's name. Let the candidate know that the interview will be recorded, for quality control and training of future interviewers.

Stick to the script as much as possible. Then you should say that you are going to start the interview. If you have questions, you should wait until the end.

Q1. How comfortable do you feel working with laptops/computers?
(1) Very comfortable, (2) A little comfortable, (3) Indifferent, (4) A little uncomfortable, (5) Very uncomfortable

Q2. Do you typically send emails or type more complex texts? Can you tell me the last time you did something like this? *OPEN ANSWER*

Q3. Have you ever used Word, Excel, or similar programs? If so, can you give me an example of something you have done with this program? *OPEN ANSWER*

Q3. Have you ever used Word, Excel, or similar programs? If so, can you give me an example of something you have done with this program? *OPEN ANSWER*
Interviewer evaluates how well the candidate did on this question, from 0 to 10

Q4. Now I will also you to do an activity. Think of a product you like and know well. It could be a type of clothing, a cell phone, a car, anything, but preferably something that you know how to describe and sell well, ok?
Can you try to convince me that I should buy this product from you or your store, instead of buying from a competitor? As if you were the seller of that product. *OPEN ANSWER*
Interviewer evaluates how well the candidate did on this question, from 0 to 10, and also writes down: (i) the product sold, (ii) the main argument, and (iii) whether it was convincing.

Q5. What would you say are your top 3 skills for a sales job, and why do you think you are good at them? It could be an example showing why you are good too. *OPEN ANSWER*
Interviewer evaluates how well the candidate did on this question, from 0 to 10

Q6. And your main disadvantages? Can you explain or give examples of how they affect you? *OPEN ANSWER*
Interviewer evaluates how well the candidate did on this question, from 0 to 10

Q7. What do you think makes you the best fit for this position, compared to your competitors? *OPEN ANSWER*
Interviewer evaluates how well the candidate did on this question, from 0 to 10

Q8. Thinking about your background and your day-to-day life, how would you say your experiences would help you to be a good fit for this position? You don't just need to give professional experiences. It could be academic, school, some leadership position, participation in social projects, volunteer work, or something else. *OPEN ANSWER*
Interviewer evaluates how well the candidate did on this question, from 0 to 10

Q9. Would you like to add any other information? *OPEN ANSWER*

Q10. [Interviewees self-administer this question on a tablet]
I see myself as a person that...

1. *Does a meticulous job*
2. *It's a little careless sometimes*
3. *It's trustworthy*
4. *Tends to be disorganized*
5. *Tends to be lazy*
6. *Perseveres until tasks are completed*
7. *Works efficiently*
8. *Make and follow plans*
9. *Is easily distracted*

Options are: (1) Totally disagree, (2) Partially disagree, (3) Neither agree nor disagree, (4) Partially agree, (5) Totally agree.

Ask if the candidate has any questions, and instruct the candidate to return to the reception for payment and final orientation.
Immediately after saying goodbye to the candidate, the interviewer responds, on a scale from 0 to 10 to each of the questions below. 0 means "Extremely bad" and 10 means "Extremely well".

1. *Overall, how well did the candidate perform?*
2. *How nervous did the candidate seem?*
3. *How focused did the candidate seem?*
4. *How professional was the candidate throughout the interview?*

Now, during the interview, the candidate... [Check all that apply]

1. *Had a shaky voice*
2. *Stuttered*
3. *Laugh nervously*
4. *Dressed in informal clothes*
5. *Used slangs*
6. *Made MANY grammar mistakes in Portuguese*
7. *Used swear words*
8. *Mentioned personal things, irrelevant to the position*
9. *Mentioned that they were religious or went to church or worship*
10. *Mentioned that they lived in a favela*
11. *Talked about where they came from (on that day of the interview)*
12. *Talked about where they lived*
13. *Talked about where they were born*
14. *Asked you personal questions*

15. *Asked you irrelevant questions for the position*

16. *Showed you know they knew something(s) about the company or the position*

17. *Used very formal language*

18. *Looked you in the eyes when answering*

19. *Avoided looking into your eyes*

20. *Was very shy*

21. *None of the above*