Expected Discrimination and Job Search*

Deivis Angeli† Ieda Matavelli‡ Fernando Secco§

Jan 28th, 2025. See the latest version here.

Abstract

We study how expected discrimination affects job applications and interview performance in three field experiments with 2,167 jobseekers living in favelas (urban slums) in Brazil. We focus on antifavela discrimination, which is overestimated by 87% of the jobseekers. Not asking for a home address only encourages white jobseekers to apply more often, likely because they can pass as nonfavela residents. Merely expecting interviewers to know one's favela address (when they actually do not) reduces average job interview performance by 0.13SD (0.3SD among white jobseekers). Hence, expected discrimination can create self-reinforcing loops and change applicant pool composition.

*We are grateful for invaluable guidance from Matt Lowe, Siwan Anderson, Jamie McCasland, and Munir Squires. Beatriz Morgado Marcoje provided unrivaled research assistance. We also thank Mackenzie Alston, Nava Ashraf, Leonardo Bursztyn, Adeline Delavande, Claudio Ferraz, Pauline Grosjean, Supreet Kaur, Ro'ee Levy, Federico Masera, Nathan Nunn, Devin Pope, Gautam Rao, Chris Roth, Rogério Santarrosa, Heather Sarsons, Colin Sullivan, Jonathan Zinman, 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) and Insper's Research Ethics Committee (Opinion N. 281/2023) and was preregistered in the AEA RCT Registry (AEARCTR-0009359).

†Global Talent Lab, deivisangeli@gmail.com

‡University of New South Wales; iedarmatavelli@gmail.com

§Analysis Group, fernandoseccoluce@gmail.com

1 Introduction

Employers often discriminate along dimensions such as race, sexual orientation, and criminal history.[1] However, jobseekers' *reactions* to expected discrimination, i.e., their expectations about (unfairly) receiving different treatment due to some characteristic they have, remain understudied. These reactions can play an important role in determining the labor market equilibrium regardless of how often employers discriminate. For instance, expected discrimination can amplify the effects of employer discrimination and create self-reinforcing loops (Coate and Loury, 1993). Jobseekers may feel discouraged from applying to jobs, especially if they overestimate discrimination (e.g., because of its salience in public debate). Expecting discrimination could also hurt interview performance, e.g., through an increase in stress, reinforcing interviewers' biases against discriminated groups. If so, policymakers may want to correct misperceptions to prevent the emergence of self-fulfilling prophecies, and firms could attract more applicants by credibly signaling procedural fairness.

This paper uses field experiments to characterize how expected discrimination affects jobseeking behaviors, from job application decisions to interview performance. We focus on the address-based discrimination perceived by favela (urban slum) residents in Rio de Janeiro, Brazil. In Rio, where 1.5 million people (22% of the population) live in favelas, address is perceived as a major source of discrimination, similar to how race is often perceived as a source of discrimination. However, unlike race, address is not generally visible, which allows us to randomize expected address discrimination by shifting its expected visibility. Specifically, we vary whether favela jobseekers may expect an HR firm (in one experiment) or an interviewer (in another) to know their address. To emulate real hiring processes, we run this HR firm ourselves, recruiting and interviewing for real full-time sales jobs. In a third experiment, we test whether a policy providing accurate information to lower jobseekers' expected discrimination effectively encourages job applications. Overall, we provide evidence that expected discrimination can affect jobseeker behaviors, in ways that can create self-reinforcing loops, and describe conditions under which different policies may be desirable.

We begin by documenting that most favela jobseekers overestimate antifavela discrimination. In a door-to-door survey, we told N=2,167 jobseekers, who would later participate in our field experiments, that we were also running an audit study. We explained that we had created fictitious résumés, randomly assigned each an address either inside or just outside a favela, and submitted those résumés to 700 sales jobs in Rio. Then, we asked the jobseekers to guess our findings, paying them based on accuracy. Over 85% predicted antifavela discrimination,

[1] See Neumark (2018), Rich (2014), and Riach and Rich (2002) for reviews of experimental evidence, and Kline et al. (2022) for a recent large-scale study documenting employer discrimination.

and approximately 60% predicted an antifavela discrimination rate (i.e., a relative callback decrease) of 50% or more. In fact, in the audit, we find (statistically indistinguishable) callback rates of 19.3% and 19.6% for favela and nonfavela résumés. While the latter does not imply that there is no antifavela discrimination at all (Kessler et al., 2019; Neumark, 2018), the gap between the predicted and observed outcomes shows that favela jobseekers are too pessimistic about antifavela discrimination, which could damage their job market prospects.

In a first experiment with jobseekers, to test whether expected discrimination affects job application decisions, we manipulated whether jobseekers could expect their address to be visible when applying for real jobs (Address Omission Experiment; N=1,303). There were three experimental arms: (i) *Address Omission*, in which an address was not needed to apply, (ii) *Status Quo*, in which jobseekers had to fill in their address on the online application form, and (iii) *Known Address*, in which true favela addresses were prefilled on the form. As the *Status Quo* arm permits address obfuscation, e.g., declaring an address in a different neighborhood or a relative's address, we included *Known Address* to prevent obfuscation, guaranteeing variation in expected address visibility.

We find no evidence that expected address visibility affects average application rates, but our (preregistered) heterogeneity tests suggest that the correlation between race and address is the main reason for the average null. Approximately 40% of the jobseekers finish the job application form (a low-cost measure of interest), and 20% attend a job interview (the "harder" outcome), with no statistically significant differences across conditions. Nevertheless, among white jobseekers (22% of our sample), interview attendance increases as we decrease expected address visibility: 13% attend in *Known Address*, 21% in *Status Quo*, and 28% in *Address Omission*, with the difference between the first and the last significant at the 1% level. This is consistent with white jobseekers expecting to pass as nonfavela residents while nonwhite jobseekers expect discrimination either way—because their race will eventually be visible or because observers will associate their race with residing in a favela. As only one-third of favela residents in Rio are white, forbidding employers from asking about home address may do little to reduce inequalities.

In a second experiment, we test whether expected discrimination can hurt interview performance (Interview Experiment; N=422). In this experiment, conducted in an office downtown, a receptionist asked jobseekers to state their name, date of birth, and address. Then, before taking the candidate to the interviewer, the receptionist stated that, to keep the process objective, "the interviewer will only know your name" (*Name-Only* condition) or "your name and address" (*Name-and-Address* condition), so the two conditions differed only by two words: "and address". To guarantee that any treatment effects came through *expected* address visibility, the

interviewers were blind to the procedure, knowing only names during all interviews.[2] After each interview, the interviewers rated the candidates on overall performance, nervousness, and professionalism, and the candidates filled out self-assessments on the same three dimensions. To maximize power and alleviate concerns about multiple hypothesis testing (MHT), we construct an index of performance for each point of view (Anderson, 2008). As our primary outcome, we average the two.

When candidates expect their address to be known (*Name-and-Address*), aggregate performance decreases by 0.13SD, compared to *Name-Only*. Unbundling the aggregate index, we find that *Name-and-Address* leads to a larger decrease in the self-assessed (0.17SD, $p<0.01$) than in the interviewer-assessed performance index (0.09SD, $p=0.28$), but we cannot reject that the effects are equal ($p=0.33$). Further evidence confirms expected discrimination as the mechanism and shows that it can indeed affect interviewer-assessed performance. For instance, we find that among those who predicted high discrimination in the audit study, *Name-and-Address* causes a statistically significant decrease of over 0.2SD on both interviewer- and self-assessed performance, which is consistent with expected discrimination leading to a self-fulfilling prophecy: If a jobseeker expects a worse evaluation due to her address, she is indeed more likely to perform worse, even when her interviewer knows only her name. Furthermore, among white jobseekers, *Name-and-Address* causes a statistically significant decrease of over 0.3SD on both interviewer- and self-assessed performance, consistent with race being the main determinant of the effect of expected address visibility, as in the previous experiment.

We ran a third experiment to understand whether a potentially more policy-relevant intervention—providing information about market-level discrimination—could encourage job applications (Information Experiment; N=690). We randomized whether surveyors told jobseekers about our audit study results, i.e., whether they revealed that we found no evidence of callback discrimination in the audit study.

Regardless of the information revealed, interview attendance was again near 20%. Together with the (mostly) null effects in the Address Omission Experiment, this result further suggests that average job application levels are not particularly elastic to expected discrimination. In light of all our experiments and data, we discuss the likeliest reasons for this low elasticity in Section 5. For instance, we have exploratory survey evidence that many jobseekers try to ignore discrimination and apply to any job for which they are a fit, perhaps as a mechanism to keep up their motivation. This mechanism would dull the effects of our treatments lowering expected discrimination, which increased application rates only when total expected discrimination could

[2]Note that our manipulation rules out self-signaling effects (whereby, e.g., reminding candidates of their address might lead them to update negatively about their skill) since all jobseekers were asked about their address before treatment.

be driven to zero (i.e., for white jobseekers under *Address Omission*).

The less pronounced treatment effects in the application stage are not at odds with the more pronounced effects on interview performance. This is because the stakes, the kinds of behaviors being controlled, and the environment are different in interviews. As interviews are high pressure and high stakes, jobseekers might have trouble dealing with expected discrimination on the spot. Accordingly, we present evidence that the effects on interview performance are best explained by "choking under pressure", which would not happen during "cold" application decisions (Baumeister, 1984; Böheim et al., 2019; Harb-Wu and Krumer, 2019; Teeselink et al., 2020). For instance, we see larger effects of *Name-and-Address* on self-reported nervousness among those who expected high discrimination at baseline or who are white, consistent with perceived discrimination leading to stress (Berger and Sarnyai, 2015; Schmader et al., 2008), and evidence that jobseekers may have a harder time controlling their behavior in the interview office. While other mechanisms, such as a reduction in perceived returns to effort, are possible, it is not obvious that they would create the observed empirical patterns.

We provide three main takeaways. First, we show that expected discrimination can exacerbate the impacts of employer discrimination, even if such discrimination is rare, and can create self-reinforcing loops by hurting interview performance. Second, our experiments shifting expected address visibility provide evidence that anonymization or blinding policies that fail to credibly conceal all correlated sources of discrimination may have limited effects or even exacerbate existing inequalities. Finally, while provision of credible information about the prevalence of market-level discrimination could prevent the adverse effects of expected discrimination, we do not find evidence that a policy providing findings of an audit study would be effective in encouraging applications.

Related literature and contributions. While many field experiments have measured discrimination in the demand side of the labor market (Neumark, 2018; Bertrand and Duflo, 2017; Rich, 2014; Riach and Rich, 2002), the supply side received little experimental attention.[3] To our knowledge, our field experiments are the first to estimate the effects of expected discrimination on jobseeker behaviors. Related field experiments randomizing either the language in job ads (Del Carpio and Fujiwara, 2023; Burn et al., 2023) or the description of the selection

[3]Some observational studies find evidence consistent with expected discrimination affecting jobseekers. For instance, people who expect discrimination may respond by generating clearer productivity signals (Lepage et al., 2022; Dickerson et al., 2022; Lang and Manove, 2011) or hide their stigma at a high cost (Agüero et al., 2023). Pager and Pedulla (2015) use administrative and survey data to show that Black jobseekers cast wider nets in their job searches and that breadth correlates with having suffered discrimination. We also build on lab studies showing that jobseekers may obfuscate their race or gender in response to expected discrimination (Kang et al., 2016; Charness et al., 2020), and there is also lab-in-the-field evidence that expected discrimination may affect on-the-job outcomes such as retention (Ruebeck, 2024) and productivity (Hoff and Pandey, 2006). See also Fryer et al. (2005).

procedure (Avery et al., 2023; Ibañez and Riener, 2018) provide suggestive but inconclusive evidence that expected discrimination can affect the job applicant mix. For instance, Del Carpio and Fujiwara (2023) find that nongendered (as opposed to gendered) job ads may encourage women applicants but may also signal better work-life balance and an inclusive culture, which can appeal differently to women. Our design lets us identify the effects of expected discrimination more sharply, for two main reasons. First, we elicit beliefs about discrimination at baseline, which allows us to estimate whether expected discrimination predicts effect sizes. Second, we designed our experiments to vary expected address visibility, keeping job desirability and other factors constant, or expected market-level discrimination, avoiding job-level confounders.

We contribute to several strands of the literature related to discrimination and labor supply. First, 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, such that statistical discrimination becomes rational (Coate and Loury, 1993; Lundberg and Startz, 1983). Glover et al. (2017) show that a similar self-fulfilling prophecy can stem solely from managers' beliefs about group differences when their initial beliefs lead them to exert less effort in supervising minority cashiers over a trial period. The lack of supervision then makes those cashiers less productive and less likely to be hired, in a self-reinforcing loop. Here, we show how expected discrimination can also generate a self-fulfilling prophecy in the matching process, exclusively through jobseekers' beliefs, as we hold the HR firm's actions constant.

Second, we contribute to the study of job interviews. We are unaware of other field experiments measuring performance in face-to-face interviews; Godlonton (2020) is likely the closest in this regard, as it experimentally measures the effects of job security on performance in training sessions. Similarly to Godlonton's (2020) finding that the effects on performance are mediated by stress, our findings help establish stress as a main driver of face-to-face performance. Regarding the connection between discrimination and interview performance, there is some indication that discrimination in interviews might be larger than on callbacks (Quillian et al., 2020; Shukla, 2024). Goldin and Rouse (2000) shows evidence that hiring of women increases after orchestras adopt "blind" auditions, also hinting at the relevance of expected discrimination: The increase could be both because evaluators lose the ability to discriminate *and* because women might perform better knowing that they will be judged based only on how they sound; our findings suggest that there is some merit to the latter explanation.[4]

Third, we contribute to a developing literature on job search. In addition to replicating find-

[4]In a small experiment with students, Word et al. (1974) articulate how even nonverbal interviewer cues triggered by a racial mismatch between interviewer and interviewees can lead to worse interview performance.

ings of excessive optimism about success rates (Mueller et al., 2021; Bandiera et al., 2023), we document pervasive pessimism about discrimination rates and study its implications. Similarly to other work on interventions improving the quality of the information available to jobseekers (Kiss et al., 2023; Belot et al., 2019), our study finds that better information may not shift average application levels, which could be due to shifts in which jobs people target or to non-monotonic responses to the relative value of applying for a job. See also Abebe et al. (2023).

Finally, our experiments randomizing expected address visibility bring the literature on stereotype threat—the idea that feeling at risk of confirming a negative stereotype may prompt a self-fulfilling prophecy (Steele and Aronson, 1995)—to the labor market. This literature overwhelmingly considers test performance or other laboratory outcomes (Spencer et al., 2016; Liu et al., 2021), and we provide field evidence of its relevance for job search.

2 Expected Discrimination in Context

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 maintain the monopoly on violence in favelas, which are home to 1.5 million people (one-fifth of the city's 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, this 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, and to self-identify as white (33% in favelas and 57% outside).

The jobseekers in our study lived in Maré, Manguinhos, or Jacarezinho, three large adjacent Rio favelas home to 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 favela and nonfavela neighborhoods (see Table A.2 for census statistics comparing these neighborhoods). We conducted most of our fieldwork in Maré, the most populous favela in Rio.

There are limited formal work opportunities in favelas. For instance, according to a census of Maré's businesses, 75% of those businesses were entirely informal and, in total, they employed only 9% of that neighborhood's working-age population (REDES, 2014). Hence, most favela jobseekers must look outside for formal jobs. Outside jobs are attractive since they provide more benefits and stability, but many favela jobseekers expect discrimination when applying to them. Formal employers typically require applicants to list a home address, meant for

assessing the applicant's distance to work. In Section 2.4, we show that many favela jobseekers believe that this address information is used to discriminate against them, regardless of distance to work.

Residents of all three of these 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 mainly an extortion racket—hold the monopoly on violence. Criminal groups were also present in the two other favelas during our fieldwork, with the police also somewhat present.[5] 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 take refuge at their homes to avoid crossfire. Workers miss work days, businesses shut their doors, and communications (internet or telephone) are hampered. It is typically unclear when a police raid ends, so the disruptions may persist for several days.

When no police raids are in progress, favela residents can typically go in and out without issues. Some may work in nonfavela 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 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 Expected Discrimination

Suppose that a group of jobseekers has a stigma, i.e., a characteristic related to (unfair) generalizations (e.g., a negative stereotype, which might bear a grain of truth; Bordalo et al. 2016) or that simply leads to antipathy in some people. If employer discrimination is differential treatment dispensed to a stigmatized group, we define *expected* discrimination as jobseekers' perceptions of the extent of such differential treatment (in relation to the treatment a similar person without that stigma would receive).

Jobseekers from favelas can expect to be the target of negative generalizations, e.g., in terms of race, income, reliability, and involvement with organized crime, and may also believe that outsiders simply dislike them, so they can expect differential treatment, as we see at the end of this section. We consider similar jobseekers who live immediately outside favelas to be the nonstigmatized reference group, so that any résumé characteristics and distance to any job are fixed. Relative to these nonfavela jobseekers, favela jobseekers might expect to receive lower callback rates, wage offers, or to be treated with more contempt. Hence, expecting discrimination can mean expecting lower returns to effort and more stress stemming from both economic

[5] See Monteiro et al. (2022), Lessing (2021), and Barnes (2022) for economic, political, and ethnographic accounts of the relationship between organized crime and the state in Rio.

hardship and feelings of injustice. We will discuss specific ways in which this can affect search behaviors throughout the paper.

2.3 Audit Study: Benchmarking Antifavela Discrimination

To benchmark current discrimination, we quantified the gap in callback rates between a favela and an adjacent nonfavela neighborhood with a new audit study. This audit study was designed to be as relevant as possible to jobseekers in our supply-side experiments, whom we would ask to predict its results.[6] We summarize the study's design and findings below; technical details and sample résumés can be found in Appendix C.

We created four fictitious worker profiles with complete high school, two male and two female. We picked random common names, which are not distinctive in terms of race, socioeconomic status, or other dimensions subject to stereotypes. Age, job experience, 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 a favela resident.

For each profile, we created two copies that differed in name, email, phone number, and address—one from Maré and one from Bonsucesso, a nonfavela neighborhood adjacent to Maré. We selected addresses that unambiguously mapped to either Maré or Bonsucesso while keeping the estimated commuting time to any job constant. Maré is a widely recognized favela in Rio, so employers can immediately distinguish the neighborhood as a favela.[7]

We collected job postings no more than two weeks old for sales positions (e.g., sales associate, telemarketing salesperson) from popular job search websites. Then, research assistants sent applications to each job posting with two different profiles with randomized addresses. We submitted 1,400 applications to 700 jobs between February and May 2023. The research assistants monitored the phone numbers and email addresses associated with each application until the end of June, coding the replies. We received 272 "callbacks," defined as invitations for interviews or on-the-job tests.

We do find no evidence of statistical differences in callback rates. The raw callback rate is 19.3% for favela and 19.6% for nonfavela, and our regression estimates do not allow us to reject

[6] Prior to our study, there was limited quantitative evidence on antifavela discrimination. Zanoni et al. (2023) found substantial antifavela discrimination in Argentina using the incentivized résumé rating method (Kessler et al., 2019). In Rio, Westphal (2014) found no antifavela discrimination (albeit with some geographical heterogeneity) using an audit study. As the Argentinian context is somewhat different and the Westphal (2014) estimates are ten years old, our audit study also makes a meaningful contribution to the measurement of antifavela discrimination.

[7] Information about the Maré-Bonsucesso callback gap is also relevant for jobseekers in the other favelas we study, since they update their beliefs about their own neighborhoods similarly to Maré residents when learning about the audit study results (see Figure A.5).

equality of means regardless of whether we add controls or job fixed effects (Table C.1). We also see no evidence that jobs in firms located in richer or more distant neighborhoods discriminate by address, and our results are similar if we recode the 27 requests for more information as callbacks.

This similarity in callback rates does not imply an absence of discrimination against favela residents. For instance, it could be the case that most discrimination occurs during interviews (Shukla, 2024) or that recruiters believe favela residents are *ceteris paribus* likelier to accept a job offer, offsetting callback differences caused by antifavela taste-based discrimination (Kessler et al., 2019). Another explanation 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 true discrimination level, it provides an objective benchmark for measuring whether jobseekers under- or overestimate antifavela discrimination.

2.4 Expected Discrimination vs. the Benchmark

In our door-to-door survey, we elicited incentivized predictions of what callback rates we would find in our audit study (see the next section for details). Before the elicitation, surveyors explained the audit study, showed the jobseeker a sample résumé, and pointed out that the address line would change from one application to another (see script in Appendix D.1.1). We focus on predictions about the jobseekers' favela of residence versus their adjacent nonfavela and compare those with the observed Maré and Bonsucesso callback rates (but reach similar conclusions if instead we focus on beliefs specifically about Maré and Bonsucesso Figure A.10).

On average, jobseekers predicted a callback rate of 63% for their adjacent nonfavela neighborhood, with 81% predicting callback rates of at least 50% (see the top panel in Figure 2). The jobseekers' guesses for favelas are closer to the audit estimates but are still too optimistic, on average: The average predicted callback rate for the jobseeker's own favela is 30%—over 50% greater than the audit study estimates.

The bottom panel in Figure 2 shows the distribution of implied discrimination rates, i.e., the percent drop in callbacks induced by the listing of a favela instead of a nonfavela address. We see that 87% predict some antifavela discrimination and 84% predict a discrimination rate larger than the upper bound of our 95% confidence interval for the discrimination rate estimated in the audit study. The median jobseeker predicts a 50% discrimination rate, almost three times larger than the audit's upper bound.

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

antifavela discrimination, since it is an incentivized prediction of an objective and relevant benchmark. We show that expected discrimination when predicting the audit study is directly relevant for the jobs in our experiments, since learning the audit results decreases how much jobseekers believe the partner HR firm would discriminate (Section 4.2).[8] As it is reasonable to assume that predicted discrimination in callback rates correlates with expected discrimination in other forms of treatment (e.g., during an interview, or on the job), and as we do not have high-quality elicitations of these other differentials, we use predicted callback discrimination as our main proxy for expected antifavela discrimination.

While the question of how beliefs about expected discrimination are formed falls beyond the scope of this paper, understanding the reasons behind expected antifavela discrimination helps us interpret the experimental results. In our survey, we asked most (N=1,497) jobseekers about the main reasons why employers may not hire people from their neighborhood. Surveyors read from a list of reasons, and respondents could agree or disagree with each. From the reason eliciting the most to that eliciting the least agreement, we have: loss of workdays because of police raids (75%); racism (68%); dislike because of cultural differences, e.g., in speech (66%); dislike of favela residents (65%); fear, e.g., of violence (60%); nepotism (57%); lower skill level (50%); difficulty in adapting to formal work (47%); and distance to work (45%). Hence, favela jobseekers have rich second-order beliefs about employers, expecting discrimination for both taste-based and statistical reasons. Notably, jobseekers seem to think that employers understand (and act on) the correlation between address and race, which might explain why white and nonwhite jobseekers react differently to shifts in expected address visibility (Section 4.4.1).

Those with higher reservation wages, who are not male, and who believe they have suffered antifavela discrimination in the past tend to predict more antifavela discrimination (Table A.3). As evidence that jobseekers understand the implications of the correlation between race and address regardless of their own race, we see that a respondent's highlighting of racism as a main reason why firms do not hire from favelas is a strong predictor of expected antifavela discrimination, but the respondent's own race is not.

Race *vs*. antifavela discrimination. As we elicited expected racial discrimination on a Likert scale (nonincentivized), we have some evidence that expected racial discrimination is even greater than expected antifavela discrimination. That is, 84% of respondents believe that firms discriminate "somewhat" or "a lot" against Black people, while 72% believe the same about people from their own neighborhood. Hence, expecting either type of discrimination

[8] Expected discrimination in the audit study also strongly correlates with other proxies of expected success probability: i) a Likert measure of antifavela discrimination and ii) a "personalized" discrimination measure comparing one's own expected future employment probability against that of a similar jobseeker in the adjacent non-favela (Figure A.11)

might already be enough to discourage many jobseekers.

Distinctive names. In our context, names are not typically perceived as distinctive in terms of race, but some names might be perceived as distinctive of lower socioeconomic status (SES). As low SES is correlated with favela residence, someone with a name that is distinctive of low SES might also expect more antifavela discrimination. The extra information contained in names is unlikely to be meaningful when employers know a lot about applicants (e.g., when reviewing complete résumés), but might be more relevant in our Interview Experiment, when name and race are the only predictors of favela residence that are always visible to interviewers.

3 Experiment Design

In early March 2023, before any randomization in the supply-side experiments, we preregistered the Address Omission Experiment and the Interview Experiment. In these experiments, to test whether expected discrimination can affect job-seeking behaviors, we randomized whether jobseekers could expect their addresses to be visible to an employer or interviewer. As our treatments typically reduce expected address visibility in relation to the status quo, one can think of these two experiments as estimating the effects of "blinding" policies on job application rates and interview performance. We launched the Information Experiment in June 2023, as we phased out the Address Omission Experiment.[9] The Information Experiment tested whether exposure to information about market-level discrimination, typically overestimated in our context, can reduce expected discrimination and encourage job applications.

3.1 Door-to-Door Survey

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.[10] It committed to giving full consideration to and fast-tracking promising applicants recruited through our pipeline. We also had the support of several NGOs in each favela, which

[9] Our introduction of the Information Experiment was the single major change to our preregistered plans. We amended our preregistration, explaining our reasoning, as we introduced this new experiment. Appendix B discusses all deviations from the initial plan and provides estimates of treatment effects on measures of application effort and other secondary outcomes listed on the preregistration, which we do not discuss in the main text for brevity.

[10] These were not perceived as strongly gendered jobs. Pilot studies did not indicate that applicants would expect significant gender bias either way; ex post, we observe that 38% of the men and 46% of the women applied to the jobs.

provided access to local networks and feedback on our procedures and research questions. Crucially, the NGOs facilitated the hiring and training of local surveyors and obtaining approval from 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. These criteria ensured most participants would be eligible for our sales jobs, but they may have excluded people who were so extremely wary of antifavela discrimination or outsiders that they had given up on looking for a formal job or responding to surveys such as ours. Hence, our design may estimate lower bounds of the effects of expected discrimination. To avoid spillover effects and maximize privacy, surveyors could interview at most one person per household, one-on-one. For the same reason, surveyors could not survey neighbors or family members of previous respondents.[11]

Survey. Surveyors invited respondents to participate in a survey about the labor market offering a participation incentive of R$5 (≈1 USD) plus a chance to win another R$500 (see Figure D.1 for photos of in-progress interviews). Then, they verified eligibility and consent and asked questions about demographics and labor market experience, without mentioning discrimination of any kind. Next, jobseekers had the option to share the information provided up to that point with a partner HR firm, which could invite them to apply for jobs if they were a fit (see the script in Appendix D.1.1). We, the researchers, operated this HR firm.

HR firm. Our choice not to present the HR firm as part of the study was deceptive to the extent that the jobseekers could not have anticipated that we would observe their interactions with the firm. This was strictly necessary for the design and was the only element of deception in this study. This separation between HR firm and academic researchers served to prevent experimenter demand effects and emulate regular labor market interactions, as research and surveys are commonly linked with local NGOs in favelas. At any rate, the HR firm invited jobseekers to apply for real jobs and indeed acted as a recruitment agent.[12]

The surveyors then told the participants that no information provided from that point onward would be shared with the HR firm, and they moved to a second block, on skills. After questions about skills and certifications, participants could take a one-minute basic algebra test paying R$0.25 for each correct answer. To build an index of skill, we aggregate i) the score in this math test ii) self-reported education level and (iii) communication skills (assessed by the surveyor

[11] These conditions were monitored, and participants who did not meet the participation criteria did not receive invitations to apply. This was feasible because we recorded precise addresses and could also check coincidences in last names before we sent job application invitations.

[12] Debriefing procedures included inviting participants who had applied for the job to a meeting to discuss the study's findings and the use of their data. For the duration of the study, we maintained a website and a contact email in case anyone searched for the HR firm.

after the survey).

The final survey block asked about job market prospects and expected discrimination. Before we launched the Information Experiment, this section included questions on *why* jobseekers believed firms would treat favela residents differently (discussed in Section 2.4).

Measuring expected discrimination. To derive our main measure of expected discrimination, introduced in Section 2.4, we incentivized jobseekers to predict the callback rates we would find in our audit study, paying an extra R$100 (≈20 USD) to the ten people whose predictions were closest to the true estimates. For both Maré and Manguinhos, we take Bonsucesso as the adjacent nonfavela neighborhood. Since Bonsucesso is not immediately adjacent for Jacarezinho, we also elicit predictions about Maria da Graça résumés from Jacarezinho residents. As our audit study covered only Maré and Bonsucesso, we elicited incentivized predictions for the other neighborhoods by initially stating that we knew the correct answer for only some of the questions.

Overview. The surveyors completed 2,392 valid surveys, which yielded 2,167 eligible participants—167 respondents did not share their data with the HR firm, and 61 of those who did provided an invalid phone number. The eligible participants were 30% men, were on average 26 years old, and were 22% white, meaning our sample is less male, is younger, and is less white than the average for favela residents (see summary statistics in Table A.1). In addition, 25% had never worked before, and 32% reported currently working full- or part-time, mostly in the informal sector.

3.2 Supply-Side Experiments

Application pipeline and internal validity. The HR firm invited all eligible participants to apply for the sales jobs, and those who applied were later invited to interview. All the eligible participants took part in the Address Omission Experiment or the Information Experiment, as indicated in Figure 1 (arrows leading from the survey stage to the application stage), with the Information Experiment introduced while we phased out the Address Omission Experiment. During the phase-out, there was an overlap of 184 participants between the two experiments. To facilitate interpretation of those two experiments, we present results for the nonoverlapping samples of those two experiments when discussing them in the main text (the results including the overlapping participants are very similar; see Appendix B).[13] All jobseekers who completed the application form and attended the interview participated in the Interview Experiment (Figure 1, arrows from the application invitations to the interview). Randomization in the interview

[13] As we launched the fieldwork in one favela at a time, the samples for each of the pre-interview studies differ in favela of origin and on some other covariates (see Table A.7).

was independent of pre-interview experimental conditions, guaranteeing internal validity; we discuss external validity in the results section.

3.2.1 Address Omission Experiment (N=1,303)

The HR firm sent personalized invitations to apply via WhatsApp to survey respondents in batches, to accommodate logistical constraints. The batch sizes varied from 50 to 117, and almost all jobseekers received their invitation up to ten days after being surveyed.

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 to apply. Those in *Status Quo* and *Known Address* received a message saying an address "is also" needed (see below). The difference between the two conditions in which an address was needed is that in *Status Quo* the jobseeker filled in the address (the usual practice in our context), which allows us to observe how often applicants obfuscate their real addresses. In *Known Address*, the form stated that the research team had already shared the jobseeker's address (along with the name and phone number) and so the applicant needed only to double-check it. Hence, *Known Address* prevented obfuscation, allowing us to test whether making an applicant's favela address fully visible affected application behavior.

WhatsApp Invitation 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! At this stage, you need to provide your **education and any courses or experience. 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.2 for the job descriptions). Next, it asked about the jobseeker's name, phone number, and address—except in *Address Omission* (Figure D.3 displays the differences across the forms). Then, it proceeded as a standard application form, asking about job experience, skills, and motivations.

Outcomes. Our main preregistered outcomes are whether the jobseeker completed the online application form and attended the job interview, which typically happened within ten days of each other. Though it is not an experimental outcome per se, we also calculate the address obfuscation rate for those in the *Status Quo* arm. We consider a jobseeker to have obfuscated her address if the declared neighborhood is neither a favela nor the postal service neighborhood of the jobseeker's real address (recorded by the surveyor after each survey).

Conceptualization. As the experimental conditions differed only in the job application procedure, it is reasonable to assume that the treatment affected only the expected value of applying to the jobs in the experiment. Then, if jobseekers assign value V, success probability p, and application cost c to our jobs, they apply when $pV - c > 0$ (with the outside option normalized to zero). We can then think of the treatment as shifting pV since the differences in application procedure were minor. For instance, in *Address Omission*, perceived pV might be larger both because a jobseeker perceives a higher success probability and because she will be less likely to suffer address-based discrimination later on the job. However, if an initial change in address visibility fails to change both the relevant success rate and expected job value, application decisions will not change.

3.2.2 Information Experiment (N=690)

The Address Omission Experiment ran until May 2023. As we phased this experiment out, we embedded the Information Experiment in our door-to-door survey to test whether provision of information about market-level discrimination, typically overestimated by favela jobseekers, can encourage job applications.

The intervention in this experiment—revealing our audit study findings to lower expected antifavela discrimination—has two potential advantages over our designs lowering expected address visibility. First, it may be more policy relevant because information provision does not rely on regulation of employers and can be transferred to any context as long as the kind of discrimination in question can be benchmarked. Second, it can provide cleaner experimental variation because it forgoes randomization in the application procedures, completely sidestepping confounders related to a change in the application procedure also changing how the HR firm is perceived (in ways unrelated to expected discrimination).[14]

[14]Such confounders are the reason why studies randomizing the language used in job ads such as Del Carpio and Fujiwara (2023) and Burn et al. (2022) cannot easily identify the effects of expected discrimination. For instance, in Del Carpio and Fujiwara (2023), gender-neutral language could imply less gender discrimination, but it also suggests better work-life balance, which can appeal differently to men and women. In this paper, similar confounders could have arisen when we randomized the application procedures in the other experiments, and that is why we kept those interventions subtle, shifting only expected address visibility.

Treatment. We randomized participants into three treatment arms: (i) *No Info*, in which no information was revealed (ii) *Favela Info*, in which they were told the favela callback rate (19.3%) after predicting the audit results, and (iii) *Full Info*, in which they were told both the favela and nonfavela callback rates (19.3% and 19.6%), revealing that we had found basically no discrimination in callback rates. As *Full Info* revealed not only the discrimination rate but also the favela callback level, we included *Favela Info* as an alternative control condition, which allowed us to hold knowledge of the favela callback level constant. See Figure 3 for the graphs the surveyors used to convey the treatment.

As in the Address Omission Experiment, the HR firm invited survey respondents to apply for our partner's jobs. There were only two differences. First, to emulate the most realistic application procedure, we used only *Status Quo* procedures, asking applicants to provide their home address. Second, since there was no randomization in the application procedure, we decreased the invitation batch sizes and made more frequent application invitations to jobseekers, one to four days after they had taken the door-to-door survey. This approach yielded a smoother flow of invitees, allowing us to allocate the HR firm's resources more efficiently, maximize our sample size, and reduce the chance that beliefs about expected discrimination would return to baseline levels.

Outcomes. In addition to the application progress outcomes used in the Address Omission Experiment, we preregistered as main outcomes the self-reported number of applications sent after two weeks (a measure we collected in an endline survey), address obfuscation, and immediate belief updates. For the latter, we 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 the elicitation by including these questions together with the items eliciting beliefs about our audit study callback rates; the surveyor introduced this set of questions with a statement clarifying that we knew the answers to only *some* of the questions.

Conceptualization. We may think of the information treatments as shifting the expected callback probability and expected discrimination for *all* sales jobs, since those were the jobs in the audit study. In this case, shifting the expected success rate p (which is downstream from the expected callback rates) can have a nonmonotonic effect on application rates. Intuitively, at a low p, an increase in p makes a marginal application much more valuable, so it is worth applying to more jobs. However, if an applicant already expects to receive "enough" callbacks, an increase in p allows her to decrease the number of costly applications while still getting enough callbacks.[15] We designed the *Favela Info* condition to fix p so that a comparison to

[15]To see this, let n be the number of applications to be submitted, p be the expected callback probability, c

Full Info could be interpreted as closing the gap in how people expect to be treated (e.g., favela residents may now expect less discrimination on the job), increasing the value of all matches after p is fixed. Obfuscation rates should increase in expected discrimination and decrease in p.

3.2.3 Interview Experiment (N=422)

The HR firm invited all jobseekers who completed the application form for a job interview at an office in downtown Rio. Attendees received a R$25 (≈5 USD) transport subsidy, enough to cover fares. We rented a reception desk and meeting rooms in a coworking space, so applicants first had to go through the building's reception and then arrive at the coworking floor. The interviews took ten to fifteen minutes each, and we scheduled them at sufficiently spaced intervals that jobseekers would rarely meet or interact on the premises. See photos in Appendix D.4.

Interview. We hired an experienced HR consultant to review our interview script and train our two interviewers. The script contained 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 (see Appendix D.4 for details). To prevent discriminatory behaviors, we instructed the interviewers to treat candidates equally and to adhere strictly to the script, never asking personal questions or about the candidates' journey to the interview location.

Treatment. We randomized expected address visibility at the job interview. A receptionist greeted the arriving candidates, asked them to state their name, date of birth, and address, and told them to wait (see Appendix D.3 for the full script). Moments later, 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 "will only know your name and address" (*Name-and-Address*). Hence, the conditions differed by two words only: "and address". The interviewers were not told about the nature of the experiment until after the end of their involvement, knowing only the participants' names prior to interviews, so that any effects on interview performance would have to have initiated with the candidate. Later, we debriefed the interviewers both to learn their impressions and to avoid participant deception (as the receptionist's statement was ambiguous about when the interviewer would learn about the addresses). Our design rules out self-signaling mechanisms (e.g., whereby a person loses confidence in her abilities when reminded of her address) since all candidates were asked to

be 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 who cares only about her first success, i.e., $V(n,p) = 1-(1-p)^n$, then we can have $V_{np}(n^*,p) > 0$ for low p and $V_{np}(n^*,p) < 0$ for high p.

state their address before treatment.

Outcomes. The interviewer evaluated the candidates immediately after each interview, and the interviewees filled out a form with self-assessment questions at the reception desk before leaving. The interviewers coded, on 0-to-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. The interviewees filled out self-assessments for the same three dimensions, so that the performance indicators totaled six. These six indicators are strongly correlated, as we should expect if they all predict real performance (Table A.8). 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 concerns about MHT, 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 (e.g., about whether she should apply for similar jobs later)—it allows us to extract a more accurate signal.

Conceptualization. The treatment shifts candidates' second-order beliefs about whether their interviewer might see them as a favela resident. Candidates might respond to this strategically or involuntarily. For instance, thinking that the interviewer knows one's address might generate feelings of unfairness due to stereotyping, leading to automatic stress responses (Berger and Sarnyai, 2015; Schmader et al., 2008). Such stress, along with the high stakes of a job interview, could lead candidates to choke under pressure, a phenomenon documented in multiple sports competitions (Böheim et al., 2019; Harb-Wu and Krumer, 2019; Teeselink et al., 2020). If a candidate simply believes her performance will be heavily discounted because of discrimination, decreasing the returns to effort, the optimal response might be to try harder to impress, which can also lead to stress, or to reduce effort if the barrier is perceived to be insurmountable. Depending on how interviewers deal with candidate reactions, an initial effect may compound or dissipate: For instance, an interviewer could try to calm down an interviewee who is too nervous, dissipating the effects of an increase in stress.

3.3 Randomization, Balance, and Estimation

Randomization for the Address Omission Experiment proceeded in batches, with strata based on baseline expected address-based discrimination and equal probability of each treatment within

and across strata. We proceeded similarly for the Interview Experiment, randomizing in batches after the jobseekers completed the application form.[16] The survey app on the surveyors' tablets randomized for the Information Experiment on the spot, also with equal probabilities. Randomizations were independent across experiments, guaranteeing internal validity.

The realized treatment assignments generated comparison groups balanced across almost all pretreatment covariates, within what would be expected from truly random assignment. In the main figures and tables, we present estimates controlling for the set of covariates for which the pairwise difference-in-means tests suggest a potentially relevant imbalance, with p<0.05, but the results are similar regardless of which controls are included (see Appendix A).[17] To plot the average application or expectation outcomes and test differences, we estimate a saturated model:

$$y_i = \alpha + \sum_{j \in T} \left[\beta^j t_i^j + \gamma^{ej} t_i^j X_i \right] + \mu X_i + \varepsilon_i \quad (1)$$

where $y_i \in \{0, 100\}$ (to yield percentages), T is either {*Status Quo*, *Address Omission*, *Known Address*} or {*No Info*, *Favela Info*, *Full Info*}, t_i^j is a dummy for assignment to arm j, and X_i is a vector of demeaned controls.[18] Thus, β^j is the covariate-corrected outcome level for outcome y in treatment arm j. For the Interview Experiment, as there are no imbalances and the outcomes are normalized, we simply estimate differences in means. In all cases, we estimate robust standard errors.

4 Results

4.1 *Address Omission Experiment*

On average, expected address visibility does not affect job application and interview show-up rates (left panel, Figure 4). If decreased expected address visibility increased expected success rates or job value, encouraging applications, *Address Omission* should have the highest application rates, and *Known Address* the lowest. Instead, we see little difference in average

[16] In these two experiments, which had $k \in \{3, 2\}$ treatments, we ordered candidates within each batch by expected discrimination in the audit and assigned the first $4k$ candidates to the one strata, the next $4k$ to another, and so on.

[17] A small number of random imbalances are expected since we could randomize only in small batches. For instance, in pairwise difference-in-means tests, we reject the null of no difference at the 5% level only five times in 102 tests (see tables A.4 to A.6). The smallest p-value among the seven omnibus F-test comparisons between each two arms within experiments is 0.1 (when we compare those in *Address Omission* to those in *Known Address*)

[18] The demeaning of covariates and the interaction $X_i t_i^j$ guarantees that we recover unconditional averages, with differences representing average treatment effects (Lin, 2013). X_i includes gender and skill in the Address Omission Experiment; residing in Manguinhos and age in the Information Experiment.

outcomes across the conditions: Application (form completion) rates hover from 41% to 45%, and interview show-up rates are just below or at 20%. The p-values for equality tests between outcomes in any two conditions are all above conventional significance thresholds. Even if we focus on jobseekers expecting high discrimination to begin with, i.e., those expecting a discrimination rate at the median of 50% or above in the audit study, we see no evidence that treatment assignment affects average application rates (right panel in Figure 4).[19]

The lack of effects cannot be explained by a failure to manipulate expected address visibility. This is because in *Status Quo*, when applicants were free to obfuscate their addresses, 45% of applicants did so, suggesting that the jobseekers indeed expected some address-based discrimination. The *Known Address* treatment was also effective in its goal of increasing expected address visibility by preventing obfuscation since only one out of the 437 jobseekers in that condition tried to provide a "corrected" address with an obfuscated neighborhood. Hence, the null result is most likely explained by expected address visibility failing to substantially shift the expected success rate or job value.

Our heterogeneity analysis suggests that expected address visibility fails to increase applications in large part because it has little or no effect on nonwhite jobseekers, who may expect expect discrimination either way. In Section 4.4, we present strong evidence that lower expected address visibility benefits white jobseekers—who can pass as nonfavela jobseekers when address is hidden—but fails to encourage nonwhite jobseekers—who can expect racial discrimination and to be stereotyped as a favela resident despite address visibility. As white jobseekers are only 22% of our whole sample, they contribute little to the average treatment effects, explaining the average null. We revisit the interpretation of the Address Omission Experiment results in Section 5 in light of all our other available evidence.

4.2 Information Experiment

First, we note that learning the callback rate observed in the audit study for *either* neighborhood shifts jobseekers' posterior beliefs about the expected callback rates that the HR firm would implement (the posterior belief) for *all* neighborhoods. For instance, we see that *Favela Info* lowers the average posterior predicted nonfavela callback rate, and that *Full Info* causes a further reduction (top left in Figure 5). This point becomes clearer when we split the results for the subsamples of those who initially over- or underestimated the favela callback rate: Posteriors about own-favela callbacks move in the same direction in which there was a surprise, as more

[19]The high-expected-discrimination group as defined here pools jobseekers who expect fairly high discrimination rates (e.g., 40%) with those who expect none. Nevertheless, the results throughout this paper are similar when we consider a cut-off of, for instance, 25%.

information is revealed (Figure A.4).[20] This cross-updating suggests that jobseekers expect a positive correlation between favela and nonfavela callback rates.

This cross-updating muddies the interpretation of *Level Info* as simply fixing the expected favela callback level, but the most dramatic reduction in expected antifavela discrimination still happens under *Full Info*. The average posterior discrimination rate for the *No Info*, *Favela Info*, and *Full Info* groups are, respectively, 35%, 28%, and 15%, with these group differences significant at the 5% or 1% level. Hence, if expected discrimination in callbacks is a major determinant of job application behavior, we should see more pronounced effects on those learning *Full Info* relative to those learning *Favela Info* (but the final effect is nevertheless a composition of effects on expected callback levels and discrimination).

Despite the belief shifts, we do not find statistically significant differences in application rates across the information conditions, even for the high-expected-discrimination group (Figure 5, bottom row). We also cannot reject the null of no effect on obfuscation rates. Only if we focus on those who initially underestimated the favela callback rate do we see that the *Full Info* treatment can decrease obfuscation, consistent with jobseekers behaving strategically (see Figure A.4, bottom left).[21]

Our endline survey, conducted over WhatsApp two weeks after baseline, generally confirms the null effects on application rates.There was no differential attrition in participation, so sample selection into the endline should not be an issue (Table 1, first column). 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; second column in Table 1), but we still find null effects on self-reported job applications after two weeks.

Evidence from two exploratory survey questions, introduced in the door-to-door survey with the Information Experiment, suggests that most of the jobseekers who were at all interested in our jobs would have applied regardless of expected discrimination. That is, (i) 68% of the respondents agreed that one should apply for all possible postings to do well in the labor market, and (ii) 84% agreed that to do well, one should not ruminate about employer discrimination.[22] If a significant share of jobseekers can successfully put discrimination out of their minds when

[20]Since the posterior prediction is about the HR firm (rather than all sales jobs), we should not expect a convergence towards the revealed callback rates. We can only predict the effects' direction, based on the weak assumption that jobseekers believe discrimination is correlated across firms.

[21]If we assume that the only way our treatments affect applications is through beliefs about callback rates, we can estimate the effect of those beliefs with two-stage least squares. This approach lets us exploit variation in how both information treatments affect beliefs about callbacks for favela and nonfavela neighborhoods but leads to similar conclusions (see Table A.12).

[22]These attitudes were elicited at the end of survey as secondary outcomes. The percentage reflects agreement on a Likert scale among those who received *No Info*. See Appendix B for treatment effect estimates on secondary outcomes.

applying, this would help rationalizing the null average effects in the Address Omission and Information experiments.

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 in response to the interviewer, who directly judges performance. This makes interview behavior very different from the "cold" decision of whether to apply.

Treatment effects. When jobseekers expect their address not to be visible, i.e., when they hear that the interviewer will know only the candidate's name, the average aggregate performance index rises by 0.13SD (Figure 6, first estimate). This index averages the interviewer's and candidate's opinions, and when we break it into its 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 the interviewer-assessed index (0.09SD) is not (p=0.28). This may suggest that the effects concentrate on self-assessments, but we cannot reject that they are equal (p=0.33). Notably, reducing expected address visibility appears beneficial across the board, as the estimates of the average effects on each of the six index components always go in the same direction (Figure 6, gray circles).

When we break down the estimates by baseline expected antifavela discrimination, we observe stronger evidence that expected discrimination acts as a self-fulfilling prophecy. For the group expecting high discrimination at baseline, we estimate statistically significant increases in performance over 0.2SD in response to *Name-Only* from both the candidate's and the interviewer's perspectives (red diamonds in Figure 6, p <0.05 for both). 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 at baseline, we cannot reject the null of no effects at the 5% level for any of the indexes or their components (orange squares, Figure 6). Hence, expecting one's address to be visible leads to worse performance among those who expect higher discrimination at baseline, even when interviewers do not have the information to discriminate. The size of the differential effect on interviewer-assessed performance among people who expect high discrimination is fairly stable (more for interviewer-assessed than self-assessed performance) even when we include other interacted covariates in the estimation (Table A.14), which suggests that other characteristics correlated with expected discrimination are not responsible for the effect heterogeneity.

External validity. The independent randomization guarantees internal validity, but since the pre-interview experiments could have changed how candidates responded to the treatment, external validity may be questioned. For instance, if the effects of *Name-Only* on interview performance were concentrated only among the few candidates who were previously told that an address was not necessary for applying, this would suggest that expected discrimination matters only in an unusual situation. To address such concerns, we test whether the treatment effects of *Name-Only* are concentrated in the subsamples with better or worse external validity. We consider "more externally valid" our results for those jobseekers whom we invited to apply through a *Status Quo* procedure and who did not learn that we found no discrimination in the audit study, while we consider the findings for the remaining candidates, who either went through an unusual application procedure or learned *Full Info*, "less externally valid."[23]

Reassuringly, the effects of *Name-Only* are not concentrated in the less externally valid case (Table 2). We estimate that the effect of *Name-Only* on aggregate performance is 0.24SD (p<0.01) for the group with more external validity and 0.03SD (p=0.7) for that with less. There is some evidence that these effects are statistically different (p=0.09), but the difference is not in the direction that would threaten external validity. When we break down interview performance into interviewers' and candidates' assessments, we see that the larger effect on the more externally valid group is due to interviewer perceptions. *Name-Only* has a strong effect on interviewer evaluations for this group (0.36SD, p<0.01) and a negative, nonsignificant effect among the subsample for which our findings are less externally valid. Since we never estimate a statistically larger effect in the latter subsample than in the former one, there is no evidence against external validity.

Effects at right tail of interviewer-assessed performance. 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 might change the composition of final hires.[24] Specifically, we can show that among those who expect high discrimination, there is a drop in the share of candidates above different cut-offs in the performance distribution. For instance, take as a benchmark the cut-off corresponding to the 90th percentile of interviewer-assessed performance index among those in *Name-Only* (mimicking a world where firms are blind to address). Among candidates expecting high discrimination at baseline, there are only half as many above that cut-off in *Name-and-Address* as in *Name-Only* (p=0.09; see Figure

[23]This strategy generates an approximately even split of participants, making the test more powerful. In Table A.9, we estimate the effects of *Name-Only* for a finer partition of the pre-interview treatments, obtaining similar results.

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

7). Put another way, among those who expect high discrimination, *Name-and-Address* ejects approximately half of the candidates from the top 10% of the distribution. No candidate who expected high discrimination in *Name-and-Address* reached the top 1%.

Mechanisms. 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 leads candidates to choke under pressure and perform worse when their addresses are visible. First, since there is little pressure when a jobseeker is deciding whether to apply, choking under pressure is consistent with our observation of the effects being more pronounced in the interview than at the job application stage.

The second empirical pattern consistent with a role for choking under pressure is how the effect heterogeneity by expected discrimination level is distributed across the six performance indicators. We see that (i) 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), (ii) there is no effect heterogeneity in the other two types of self-assessed perceptions, and (iii) the effects on all three components of the interviewer-assessed index are concentrated in the high-expected-discrimination group. These suggest that an increase in the expected discrimination leads to a stress response, which is then reflected in the interviewer's assessesments.[25] While not ruled out, a mechanism whereby discrimination reduces the perceived returns to effort during interviews would not necessarily generate such increases in nervousness when addresses are visible in the high-expected discrimination group. A similar argument applies to motivated reasoning.

Third, we see some evidence that jobseekers might find it difficult to be strategic at the office, 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 in their addresses on the online form) and made it to the interview, the same jobseekers are 20% (5.7 p.p.) less likely to obfuscate their addresses at the interview office ($p<0.01$). In addition, the performance indicator most affected (on average) by *Name-Only* is the self-perceived professional behavior, suggesting that jobseekers do not regulate their behaviors as much when they believe their address 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 considering their behavior *less* professional when their addresses are visible.

[25] Note that the effect on nervousness could go in the opposite direction if *Name-Only* led to a higher cognitive load among those who try to pass as nonfavela residents. That is, the pressure to be careful and not reveal any hints one is actually from the favela could overpower the nervousness induced by expecting discrimination, but we see the opposite.

4.4 Heterogeneity Analysis

We preregistered four heterogeneity tests: by expected discrimination, race, skill, and gender. The heterogeneity by expected discrimination is key to confirming our mechanism of interest, the race heterogeneity allows us to observe how correlated sources of discrimination interact, the skill heterogeneity could tell us how expected discrimination changes the talent pool available to employers, and the gender heterogeneity could tell us whether favela men, who are likelier to be gang members, react more to expected discrimination. In this section, we take a systematic approach to rigorously test for effect heterogeneity across these dimensions.

As our research design included three experiments, multiple treatments, and multiple outcomes, testing for heterogeneity along the four dimensions above in each single case can lead to false positives due to MHT. For completeness, we present "naive" effect heterogeneity estimates in Appendix Tables A.6 to A.9, but our main approach to interpret effect heterogeneity avoiding MHT issues is to i) reduce the number of tests, by focusing on main outcomes and simultaneously testing for effect heterogeneity across all experiments, and ii) correct for the fact we test four heterogeneity hypotheses. We operationalize this strategy by stacking the data from the three experiments and running the following regression:

$$y_{ie} = \alpha^e + \sum_{j \in T} \left[\beta^{ej} t_{ie}^j + \gamma^{ej} t_{ie}^j h_i \right] + \nu^e h_i + \mu^e X_i + \varepsilon_{ie} \tag{2}$$

where y_{ie} is the outcome of individual i in experiment e, α^e is an experiment-specific intercept, $T = \{Address\ Omission, Known\ Address, Favela\ Info, Full\ Info, Name\text{-}Only\}$, t_{ie}^j is a dummy for whether individual i received treatment j in experiment e, h_i is a dummy that depends on the dimension of heterogeneity we are testing (expecting at- or above-median discrimination at baseline, being a white jobseeker, being high-skill, or male), and $\mu^e X_i$ allows us to control for potential imbalances (Section 3.3).[26] The estimate of γ^{ej} is then numerically identical to when we estimate the effect heterogeneity one experiment at a time and, after estimating the variance-covariance matrix with individual-level clusters, we conduct an F-test for whether any γ^{ej} is nonzero.

Our estimation of Equation 2 yields five heterogeneity coefficients for each of the four dimensions (20 total), which obviates the MHT issues involved with trying to interpret each co-

[26] To test whether each covariate predicts treatment effects over and above the other covariates, one could add other interacted covariates to equation 2. For completeness, we show these results in Table A.11. The point estimates are similar, but they should be interpreted with caution since the additional covariates may be "bad controls." To see this, consider that race could have a causal effect on skill level, expected discrimination, or employment status, so the predictive power of race could be diluted not because these other characteristics cause the heterogeneity but because they are simply correlated with race.

efficient on its own (Table 3). Across the experiments, we can reject the null of no effect heterogeneity only for race (p=0.01 in the F-test, or p=0.04 after a Bonferroni correction for MHT). Hence, we can be confident that race predicts some of the effects and that interpreting its role can improve our understanding of mechanisms, sharpening our derivation of policy implications. Notably, we cannot reject the null that the baseline expected discrimination level does not predict treatment effects across experiments. That is likely because this test has less power to reject the null if expected discrimination has a first-order effect on behaviors only in "hot" (interview) conditions.

Next, we analyze the role of race first in the experiments shifting expected address visibility and then in the Information Experiment, since the different types of treatment allow for different mechanisms.

4.4.1 Race and Address Visibility

The role of race in the interview experiment is the simplest to rationalize since race was one of two predictors of favela residence that was always visible to interviewers (the other predictor was the candidate's name, which might predict SES). Other cues were rare in our interviews: only 4% of the candidates mentioned living in a favela, and only 8% deviated from formal language norms. Other (potential) predictors of address, such as education level or musical taste, would not be visible during the interviews by design, as the interviewers closely followed a script. Finally, when we debriefed the interviewers, they told us that while they had noticed that a substantial share of the candidates had low SES, they could not guess all interviewees were favela residents.

The fact that race is a key predictor of address in our interviews has consequences, since candidates understand that firms act on the correlation between address and race (Section 2.1). For white candidates, being told that their interviewer would know only their name enables them to avoid all kinds of discrimination, and we see that this is meaningful for interview performance: *Name-Only* increases white candidates performance indexes by over 0.3SD from both the candidate's and interviewer's perspective (Figure A.9). Nonwhite candidates might expect substantial discrimination despite what they are told because they may always expect i) racial discrimination (which is substantial; see Section 2.4) and ii) to be stereotyped by interviewers, who would think of them as favela residents even under *Name-Only*.[27] Consistent with this

[27] A fully Bayesian interviewer who knows only the candidate's race would guess that 13% of white and 23% of nonwhite candidates are from favelas, so nonwhite jobseekers should expect twice the address-based discrimination in that case. The 23% posterior alone could be enough to create discrimination, but we should also expect nonwhite jobseekers to be stereotyped as favela residents, as it is common in the context. See Bordalo et al. (2016) for theory and evidence of stereotyping (exaggerating differences between groups) in general contexts.

picture, the effect of *Name-Only* on aggregate performance is 0.24SD smaller among nonwhite candidates (Table 3). Breaking down this effect differential by the six components of performance, we see a pattern similar to what emerged when we discussed heterogeneity by baseline expected discrimination, which suggests that the mechanism at play is also a kind of choking under pressure when white jobseekers believe their addresses are known (Figure A.9).[28]

The same rationale we used to understand race heterogeneity in the interview experiment can explain the results of the Address Omission Experiment as long as jobseekers expected that their race would eventually be visible to employers. Then, white jobseekers should expect higher success probabilities when we lower their expected address visibility, meaning that those who did not need to provide an address (*Address Omission* arm) should expect a higher success rate than those who did (*Status Quo*) and the latter should expect a higher success rate than those whose favela addresses were prefilled (*Known Address*). Interview attendance should increase in that success probability, and it does so for white jobseekers: It is 14% in *Known Address*, 23% in *Status Quo*, and 29% in *Address Omission*. Nonwhite candidates attend interviews at approximately the same rate (17% to 19%) across arms, consistent with their expecting their race to eventually become visible, which would negate the effects of expected address visibility.[29]

Since i) the same causal mechanism can coherently explain the effect heterogeneity by race across the experiments shifting expected address visibility and ii) the joint hypothesis test shows that the existence of heterogeneity by race specifically in the Address Omission and Interview experiments is robust to MHT concerns (p=0.023, or p=0.092 after a Bonferroni correction), we are confident in discussing the derived policy implications in Section 5.

4.4.2 Race in the Information Experiment

As the information that we found no antifavela discrimination in callbacks lowers expected market-level discrimination, without a direct effect on expected address visibility, the mechanisms behind the effect heterogeneity by race in this experiment might differ from those previously discussed. For instance, learning that we found no discrimination in the audit can encourage nonwhite more than white candidates: As address and race are correlated, learning

[28] As name was the other predictor of address that was always visible to interviewers, the same mechanism suggests that effects might concentrate on white jobseekers with names that are not distinctive of low SES. An exploratory heterogeneity analysis, after classifying interviewee names into those that are more or less likely to be perceived as distinctive of low SES, confirms this prediction (Table A.15).

[29] This pattern is less clear-cut if we look at application instead of show-up rates (Figure A.6), but this does not threaten our interpretation for two main reasons. First, race is still very predictive of the effect of *Address Omission* against the other conditions pooled; i.e., estimating $application_i = \alpha + \beta \times Address\ Omission_i + \nu white_i + \gamma Address\ Omission_i white_i + \varepsilon_i$ yields $\hat{\gamma}$=16.8 p.p. (p=0.02). Second, the apparent decrease in applications in response to address visibility among nonwhites is probably a false positive since it does not have a counterpart in show-up rates or in the Interview Experiment.

that firms do not use address information to discriminate implies they also do not use that information for racial discrimination, meaning that racial discrimination is also low, which should be more encouraging to nonwhite candidates. Learning the audit results could also increase perceived labor market competition: If white jobseekers become convinced that total discrimination is low across the job market, this means that they will compete with many more nonwhite candidates for each job. Hence, testing for effect heterogeneity in the information experiment probes the relevance of intrinsically different mechanisms, independent of the mechanisms involved in the interaction of race and expected address visibility.

There may be some truth to those narratives, but looking at all the data we have available suggests even more nuance. The negative coefficient on *Full Info*×white in Table 3 suggests that nonwhite jobseekers are relatively encouraged to apply to more jobs after learning the audit results. There is also an apparent net discouragement effect of *Full Info* on white jobseeker application behavior (Figure A.7), consistent with increased perceived competition. On the other hand, we see *Full Info* being relatively more encouraging for white jobseekers when the outcome is the intended number of applications immediately after treatment (a secondary outcome; Table B.4) and no heterogeneity by race in the effect of *Full Info* on the number of job applications declared in our endline survey (Table A.10). Hence, a better explanation could be that white jobseekers are encouraged by *Full Info* but that it also leads them to shift their focus to nonsales jobs, where they perhaps expect less competition. This shift in which jobs white jobseekers apply for is consistent with work in the job search literature finding that provision of information that could increase total search effort may instead lead to adjustments on other margins, such as which jobs participants apply for (e.g., Kiss et al. 2023; Belot et al. 2019).

Nevertheless, we are cautious about overinterpreting the observed patterns, since i) we do not have data on exactly what other jobs white jobseekers applied for and ii) the stand-alone test of whether race predicts the effects of the information treatment on show-up yields p=0.20 after a Bonferroni correction (Table 3). The main takeaway from the Information Experiment still is that information about market-level discrimination is not an effective tool to encourage more search effort.

5 Discussion

5.1 Revisiting the Experimental Results

We can bring together the evidence from the different parts of Section 4 to strengthen our understanding of the effects of expected discrimination.

Expected discrimination affects interviewer-assessed performance. While the estimated average treatment effect of *Name-Only* on the interviewer-assessed performance index (0.09 SD) is not statistically significant, three features of our body of evidence support the hypothesis that expected discrimination indeed affects interviewer-assessed performance. First, we estimate statistically significant effects of *Name-Only* on the subgroups for whom expected address visibility would matter the most: those who expect high address-based discrimination at baseline (0.22SD effect) and those who are white (0.31SD). Second, the same choking-under-pressure pattern helps explain the results for both white and high-expected-discrimination candidates, reinforcing the same causal story. Third, we also estimate a statistically significant increase in interviewer-assessed performance (0.36SD) in the subsample of candidates for which our results are more externally valid, which should receive a higher weight when considering policy implications.

Null *average* effects on job applications. What is most likely to explain the null average effects on job application decisions, and how does this result fit with the more noticeable treatment effects on interview performance? The main explanation for the null effects should not be that expected discrimination never matters for application decisions since white jobseeker interview attendance decreases with higher expected address visibility, consistent with white jobseekers also being the most affected in the Interview Experiment. Rather, we should consider explanations about why expected address visibility and market-level information may fail to change behavior. As we see no effects of shifting address visibility among those who expected high antifavela discrimination at baseline, it is possible that many jobseekers successfully abstract from discrimination in "cold" situations, as indicated in our exploratory survey questions. Alternatively, jobseekers might simply expect less discrimination conditional on their receiving a personalized application invitation, dulling effects on application outcomes.

The same mechanisms might yield the average null effects of information provision, and others are also possible. For instance, we cannot rule out nonmonotonic effects of expected success probability since jobseekers cross-update between favela and nonfavela callback rates so *Favela Info* does not hold beliefs about callback levels perfectly constant. Another possibility is that learning that there is no discrimination in callbacks also increases perceived job market competition among white jobseekers, attenuating the effects. Finally, jobseekers may not find market-level statistics persuasive enough to translate their updated beliefs into action (i.e., the intervention generates weak posteriors) or may just be uncertain about how to do so.[30] At any

[30] We collected a Likert-scale measure of confidence in posterior beliefs for some participants of the Information Treatment, as an (non-preregistered) check. Somewhat consistent with the idea of weak posteriors, only about one-third of respondents who learned the full audit results stated that they were "very" or "extremely" confident about their beliefs.

rate, none of those explanations are at odds with a choking-under-pressure mechanism leading to worse interview performance or with the correlation between race and address being the main determinant behind the effect of expected address visibility.

5.2 Remarks and Implications

Self-assessed performance. Even if the effects of expected address visibility were restricted to self-assessed performance, 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. In addition, note that, in "regular" interviews, discriminatory behavior among interviewers can exacerbate any effects of expected discrimination while, in our interviews, this channel was shut down since the interviewers knew only names and stayed on script. Finally, beyond any behavioral effects, expected discrimination can undermine jobseekers' psychological welfare (Pascoe and Smart Richman, 2009; Schmitt et al., 2014), as we show that it leads to negative interview experiences.

Anonymization and blinding. Our experiments have implications for policies that restrict the information recruiters may access. First, consider policies that reduce the visibility of a stigma at the callback stage, such as résumé anonymization or bans on requests for some specific information. Our results suggest that we should not expect such policies to change applicant behavior across the board, especially if most jobseekers bear an always-visible trait that is strongly correlated with the trait(s) hidden by the policy. Such policies might encourage applications only for groups that can continue to hide their stigmas later on, as was the case for white jobseekers in our experiments randomizing expected address visibility. 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 that these policies should be treated with even more caution.

On the other hand, there are reasons for optimism 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, so if jobseekers are made aware of credible blinding or antidiscrimination policies, those who expect discrimination might then perform better in interviews. Policies that do not allow for correlated sources of discrimination to inform each other (e.g., audio-only, text-only, or metaverse interviews) should also be more effective since jobseekers can anticipate interviewers' reactions to the correlated signals; AI-intermediated candidate selection is also a promising alternative, as shown in Avery et al. (2023).

Intersectionality. Our results show that interventions ignoring race as a correlated source

of discrimination can have heterogeneous and unintended results, such as increasing racial inequalities. This finding resonates with the idea of intersectionality, i.e., that overlapping stigmas interact in ways simple additive effects cannot summarize, and that first-best policies take those complex interactions into account (Carvalho et al., 2022; Crenshaw, 1989). Nevertheless, our findings also suggest that even when we know two stigmas might interact, it may be hard to predict policy consequences. For instance, a policymaker might expect that hiding addresses in interviews would benefit nonwhite more than white jobseekers, which would be the case if jobseekers 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 mostly to white jobseekers, which might run against the intuitions of many. Hence, an information-constrained policymaker could justifiably experiment with race-blind policies to address expected antifavela discrimination as a first-order approximation and then iterate based on its results.

6 Conclusion

This paper provides evidence that expected discrimination can affect jobseekers regardless of the employer discrimination level and work as a self-fulfilling prophecy in job interviews, potentially contributing to the labor market inequalities observed in administrative data. We document a novel mechanism through which expected discrimination can act as a self-fulfilling prophecy: by hurting interview performance, instead of lowering returns to human capital investment (Coate and Loury, 1993) or interacting with actual on-the-job discrimination (Glover et al., 2017). In addition, race heterogeneity in treatment effects suggests that policymakers should be mindful of intersectionalities.

This paper joins recent work showing that discrimination tends to be overestimated (Haaland and Roth, 2023; Aksoy et al., 2023; Angeli and Lowe, 2024), raising questions about what can shape beliefs about expected discrimination. A general bias toward overestimation is concerning not only because of the effects discussed in this paper but also because, if employers engage in statistical discrimination in response (or due to inaccurate beliefs, as in Bohren et al. 2023), then those initially inaccurate beliefs may become true. Such a perverse equilibrium could be avoided if a policymaker could simply convey truthful market-level information, but our information experiment shows that this is not a simple task. Future research on the determinants of discrimination-related beliefs could point to better solutions. Negativity bias in the demand for news supplied by media or peers (e.g., Robertson et al. 2023) may play an important role in sustaining these misperceptions. As news stories often home in on compelling narratives,

market-level statistics may not be persuasive enough to change high-stakes actions (Graeber et al., 2024), and alternative policy tools may be necessary.

Future research identifying precisely why expected discrimination is (more) relevant at the interview than at the application stage can also sharpen our policy implications. Our results suggest that choking under pressure might be behind the negative effects on interview performance, so that experiments varying pressure, or whether the interview has a face-to-face element, could shed some light on mechanisms. If choking under pressure is indeed to blame, interventions focused on controlling behavior and decreasing anxiety, such as cognitive behavioral therapy, could counter its effects. Alternatively, improved access to skill signaling technologies (Abebe et al., 2021; Carranza et al., 2020) could also decrease the relevance of interviews.

Finally, 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 lower the discrimination jobseekers expect from those firms. If DEI commitments indeed remove a handicap faced by jobseekers who anticipate discrimination and help recruiters identify talent, 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.

_ , Stefano A Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, Simon Quinn, and Forhad Jahan Shilpi, "Matching frictions and distorted beliefs: Evidence from a job fair experiment," Technical Report, working paper 2023.

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.

Angeli, Deivis and Matt Lowe, "Do Virtue Signals Signal Virtue?," *Working paper*, 2024.

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.

Bandiera, Oriana, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali, "The search for good jobs: evidence from a six-year field experiment in Uganda," Technical Report, National Bureau of Economic Research 2023.

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

Baumeister, Roy F, "Choking under pressure: self-consciousness and paradoxical effects of incentives on skillful performance.," *Journal of personality and social psychology*, 1984, *46* (3), 610.

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

Belot, Michele, Philipp Kircher, and Paul Muller, "Providing advice to jobseekers at low cost: An experimental study on online advice," *The review of economic studies*, 2019, *86* (4), 1411–1447.

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

Bertrand, Marianne and Esther Duflo, "Field experiments on discrimination," *Handbook of economic field experiments*, 2017, *1*, 309–393.

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.

Bohren, J Aislinn, Kareem Haggag, Alex Imas, and Devin G Pope, "Inaccurate statistical discrimination: An identification problem," *Review of Economics and Statistics*, 2023, pp. 1–45.

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.

Godlonton, Susan, "Employment Risk and Job-Seeker Performance," *Journal of Human Resources*, 2020, *55* (1), 194–239.

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

Graeber, Thomas, Christopher Roth, and Florian Zimmermann, "Stories, statistics, and memory," *The Quarterly Journal of Economics*, 2024, p. qjae020.

Haaland, Ingar and Christopher Roth, "Beliefs about racial discrimination and support for pro-black policies," *Review of Economics and Statistics*, 2023, *105* (1), 40–53.

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.

Ibañez, Marcela and Gerhard Riener, "Sorting through affirmative action: Three field experiments in Colombia," *Journal of Labor Economics*, 2018, *36* (2), 437–478.

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.

Kiss, Andrea, Robert Garlick, Kate Orkin, and Lukas Hensel, "Jobseekers' beliefs about comparative advantage and (mis) directed search," *Available at SSRN 4593303*, 2023.

Kline, Patrick, Evan K Rose, and Christopher R Walters, "Systemic discrimination among large US employers," *The Quarterly Journal of Economics*, 2022, *137* (4), 1963–2036.

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.

Lin, Winston, "Agnostic notes on regression adjustments to experimental data: Reexamining Freedman's critique," 2013.

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.

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.

Mueller, Andreas I, Johannes Spinnewijn, and Giorgio Topa, "Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias," *American Economic Review*, 2021, *111* (1), 324–363.

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.

Robertson, Claire E, Nicolas Pröllochs, Kaoru Schwarzenegger, Philip Pärnamets, Jay J Van Bavel, and Stefan Feuerriegel, "Negativity drives online news consumption," *Nature Human Behaviour*, 2023, *7* (5), 812–822.

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

Schmader, Toni, Michael Johns, and Chad Forbes, "An Integrated Process Model of Stereotype Threat Effects on Performance.," *Psychological Review*, 2008, *115* (2), 336.

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.

Shukla, Soumitra, "Making the Elite: Coded Discrimination at Top Firms," 2024.

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.

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 nonfavela 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: Interview Treatment Effects by Treatment Conditions Before Interview

	(1) Aggregate performance index	(2) Interviewer-assessed performance index	(3) Self-assessed performance index
Name-Only × *Status Quo* $\times non-FullInfo$	0.23***	0.35***	0.12
	(0.09)	(0.12)	(0.10)
Name-Only × *All other pre-interview conditions*	0.03	-0.13	0.19**
	(0.08)	(0.11)	(0.08)
Observations	422	422	422
P-value for same effect on more vs. less externally-valid conditions	0.09	0.00	0.56

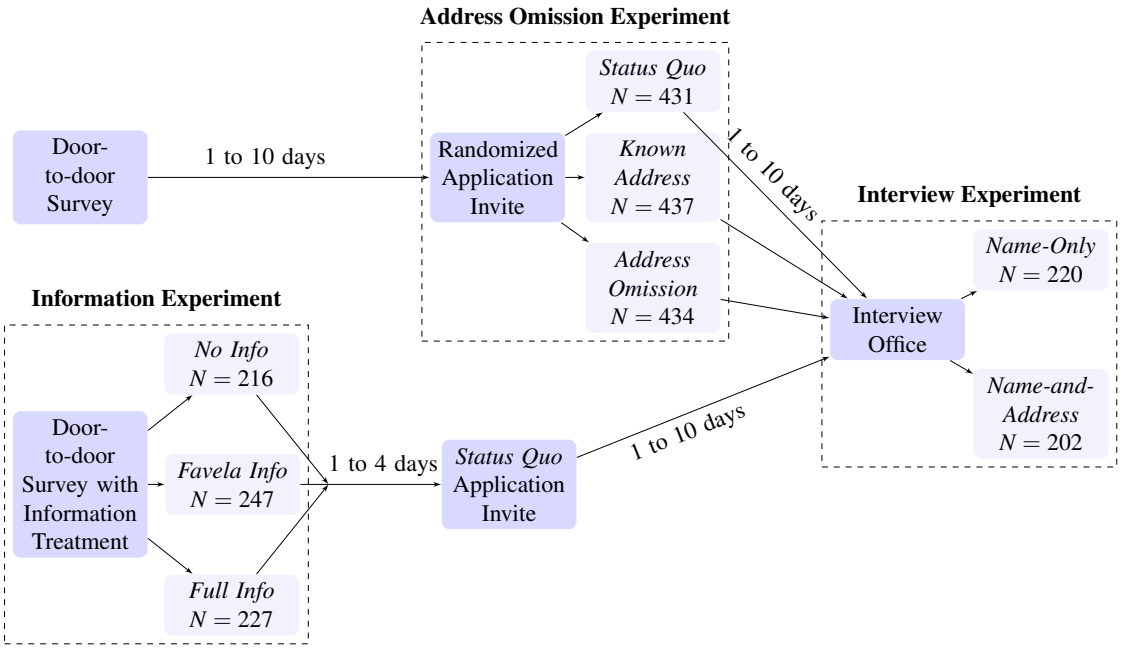
Note: OLS estimates for the effects of *Name-Only* on interview performance indexes for the groups with better (who went through *Status Quo* applications and did not learn the full audit results) and worse external validity (all others). Regressions fully control for treatment assignment in previous experiments. The last row compares the two regression coefficients displayed in each column. Robust standard errors between parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 3: Across Experiments, Only the Heterogeneity by Race is Robust

	Experiment-specific outcomes			
	(1)	(2)	(3)	(4)
Agg. interv. performance (SD) *Name-Only* $\times h_i$	0.24*	0.24**	0.07	−0.12
	(0.14)	(0.12)	(0.13)	(0.12)
Show-up (%) *Address Omission* $\times h_i$	7.27	−3.90	−2.18	−1.16
	(6.74)	(5.56)	(5.61)	(5.43)
Show-up (%) *Known Address* $\times h_i$	−8.23	−5.61	5.14	−6.83
	(6.29)	(5.45)	(5.80)	(5.28)
Show-up (%) *Full Info* $\times h_i$	−18.00**	−5.99	−2.89	11.80
	(9.02)	(7.66)	(7.95)	(7.64)
Show-up (%) *Favela Info* $\times h_i$	−0.23	2.61	0.29	14.37*
	(9.92)	(7.71)	(8.05)	(7.62)
Heterogeneity variable h	White	High $\mathbb{E}$[disc]	Male	High Skill
Any heterogeneity by h, p-value	0.01	0.242	0.827	0.191
Heterogeneity by h in $\mathbb{E}$[address visib.] treatments, p-value	0.023	0.143	0.575	0.364
Heterogeneity by h in information treatment, p-value	0.049	0.502	0.901	0.136
Clusters	2032	2032	2032	2032
Observations	2,415	2,415	2,415	2,415

Note: This table presents a summary of the heterogeneity in treatment effects across experiments and tests if we can reject the null of no heterogeneity across experiments accounting for multiple hypothesis testing. Each column presents the regression coefficients, from a stacked regression, on the interaction of the experiment-outcome, treatment and h (equation 2), which are numerically the same as outcome-specific regressions.

Figure 1: Experimental Design



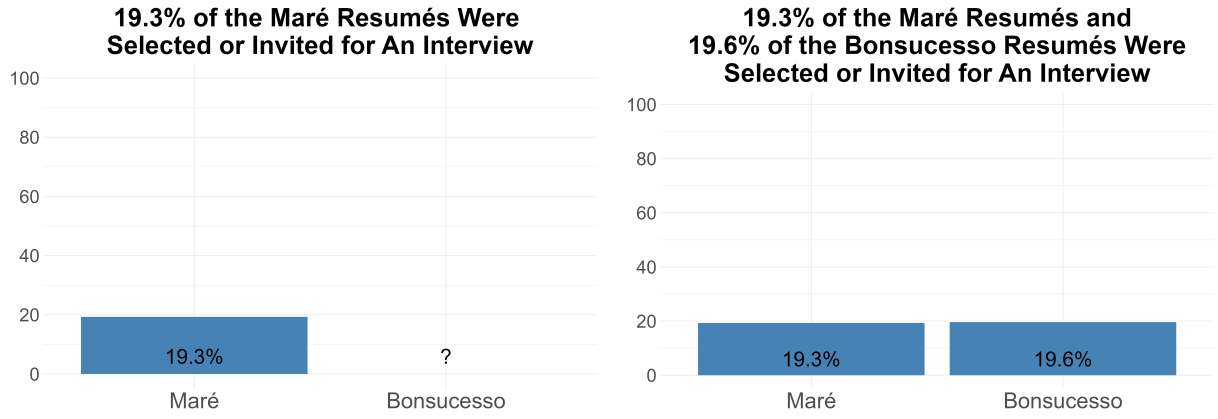
Note: The figure shows how the three supply-side experiments fit together. Arrows indicate the flow of jobseekers who answer a door-to-door survey, receive an invite to apply, and can also interview for a set of real jobs. The upper pipeline, which includes the address omission Experiment, happened from March to May 2023. The bottom pipeline was included later. Time intervals over arrows exclude outliers. Dashed boxes delineate experiments. Lighter-blue nodes are randomized experimental conditions, each with its number of participants. 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 nonfavela 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 the discrimination estimated in the audit study.

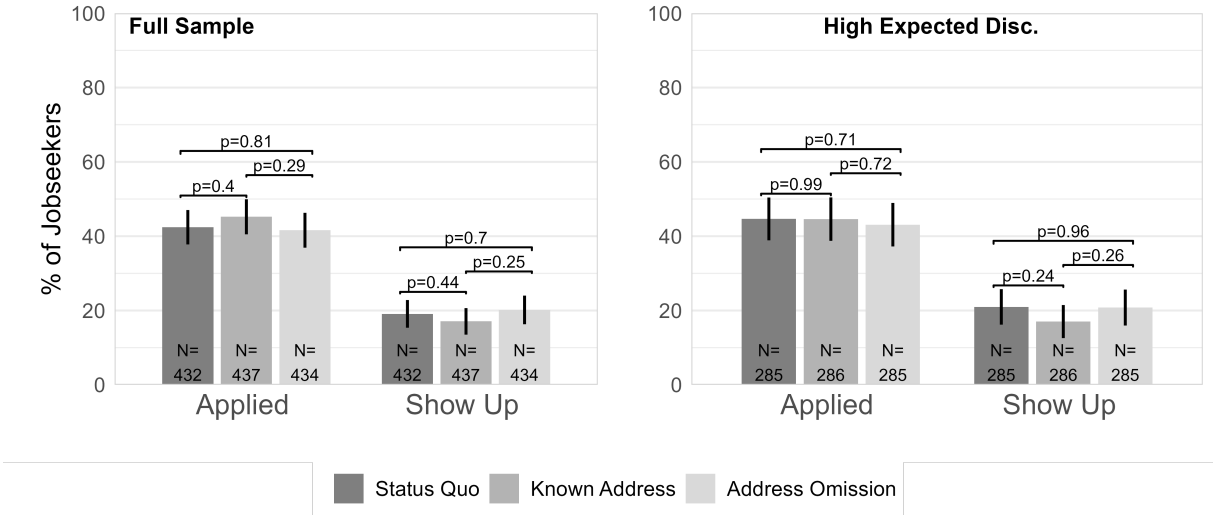
Figure 3: Information Treatment Delivery



(a) *Favela Info* (b) *Full Info*

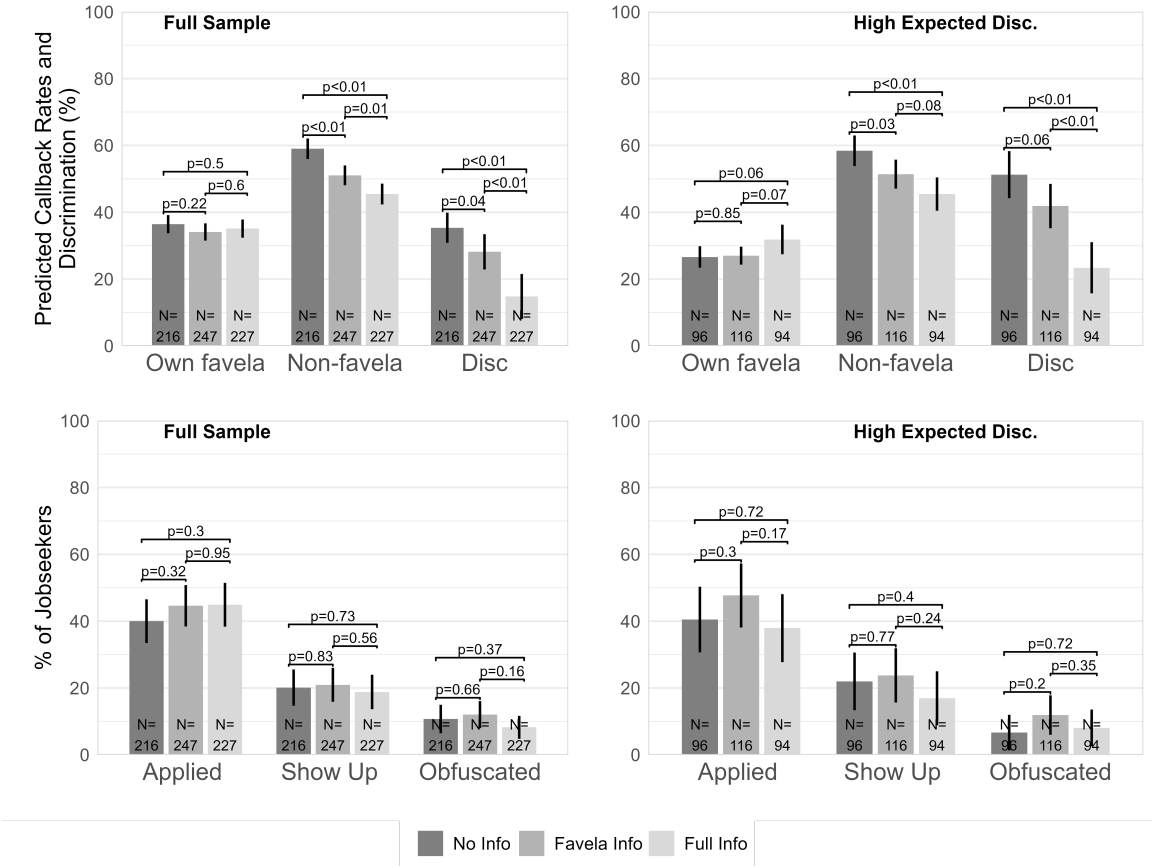
Note: This Figure shows the images we used to convey the information experiment. We showed either one of the plots (or none) to participants immediately after predicting the audit study results. The surveyor read the text above each graph when showing it to the respondent.

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



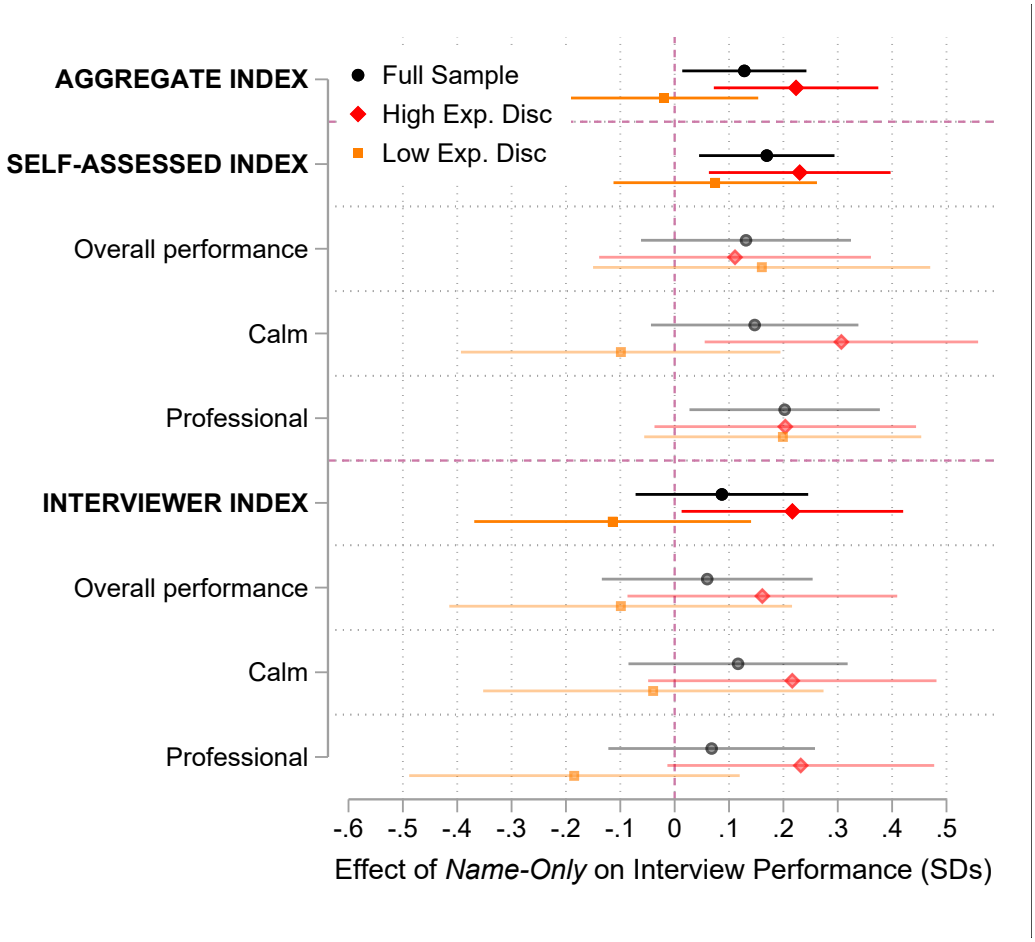
Note: This figure displays shares of all jobseekers in the address omission Experiment reaching each stage of the application process. Applied means finishing the online application form, and Show Up means showing up at the interview. The left panel shows results for the full sample, and the right panel 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



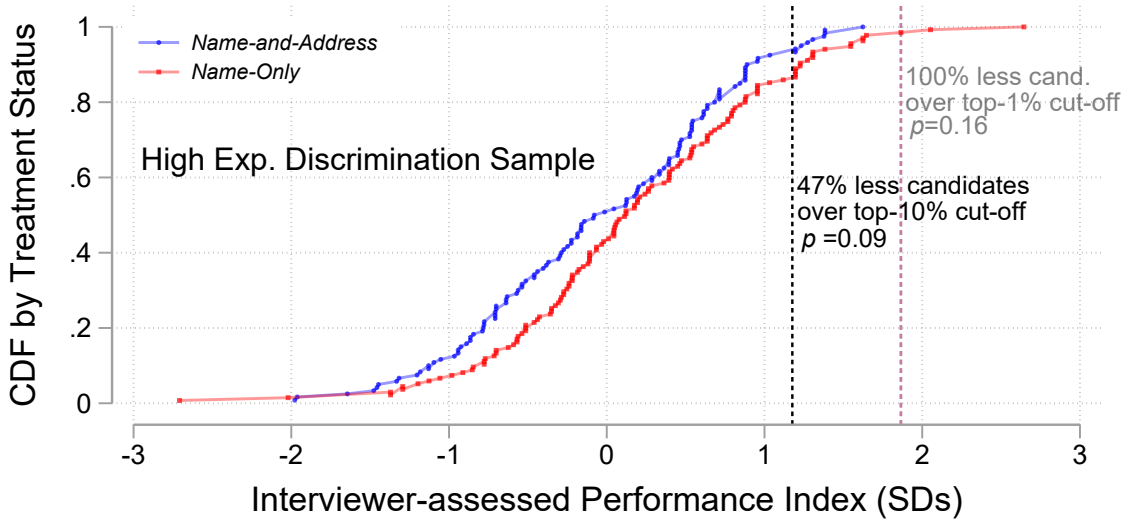
Note: The top row of graphs displays average posterior beliefs of what callback rates the HR firm would implement for jobseekers in each experimental condition. Nonfavela and Own favela stands for the callback rate prediction for a respondent's favela and adjacent nonfavela. Disc is the implied percent drop in callback rate due to the favela address. The bottom row displays outcomes from the application process. Applied means finishing the online application form, Show Up means attending the interview, and Obfuscated means declaring 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 3.2.3. 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 Address Visibility Decreases the Share of Individuals in the Right Tail of Interviewer-assessed Performance Among Those Who Expect High Antifavela 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

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	26.90	6.25	19	42	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 searched in the last 7 days	0.49	0.50	0	1	2,167
Surveyor-assessed comm skills (Likert scale, 0-5)	2.79	1.11	0	4	2,158
Math test score	6.96	2.55	0	17	2,081
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
Reservation wage (USD)	251.75	106.87	-20	2,200	2,166
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
Nepotism (is reason, 0/1)	0.57	0.50	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. Age was calculated based on the declared date of birth. Race, gender, education, and work experience were declared. "Actively searched last in the last 7 days" refers to taking any specific action to find a job (e.g., submitting a résumé) in the last seven days. "Surveyor assessed comm skills" comes from Likert-scale questions about how easily the jobseeker understood and answered the survey. Math test score is the number of multiple-choice math questions answered correctly within a minute during the survey. Reservation wage was elicited by asking for the lowest wage for which a person would accept a full-time job in their area of expertise in Downtown Rio. The last eight variables are dummies for whether the jobseekers agreed a specific reason was important for explaining why jobseekers might avoid hiring from favelas. Those nine questions were removed after we introduce the information experiment, to control survey duration.

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 for relevant neighborhoods in Rio. Bonsucesso was the adjacent nonfavela for surveys in Maré and Manguinhos. Maria da Graça was the adjacent nonfavela for Jacarezinho.

Table A.3: Expected Discrimination Predictors

	(1) Expects > 50% disc in audit	(2) Expects > 50% disc in audit
Age	0.001	0.000
	(0.002)	(0.002)
Male (0/1)	-0.070***	-0.042
	(0.024)	(0.028)
White jobseeker (0/1)	-0.026	-0.042
	(0.026)	(0.030)
Some college (0/1)	0.098**	0.055
	(0.040)	(0.049)
Completed regular high-school (0/1)	0.042	0.013
	(0.032)	(0.036)
Working now (0/1)	0.019	-0.034
	(0.029)	(0.034)
Holds a formal job (0/1)	-0.048	-0.011
	(0.039)	(0.046)
Ever worked (0/1)	-0.019	0.017
	(0.028)	(0.034)
Actively searched last week (0/1)	0.018	0.017
	(0.022)	(0.026)
High skill (0/1)	-0.025	-0.015
	(0.025)	(0.030)
Heard of people refused job/fired due to address (0/1)	0.028	0.045
	(0.026)	(0.030)
Believes has been refused job/fired due to address (0/1)	0.106***	0.078***
	(0.027)	(0.030)
Reservation wage (USD)	0.000**	0.000
	(0.000)	(0.000)
Distance to work (is reason, 0/1)		0.002
		(0.026)
Missing days because of police raids (is reason, 0/1)		0.049
		(0.031)
Lower skill (is reason, 0/1)		-0.001
		(0.028)
Difficulty adapting to work (is reason, 0/1)		-0.021
		(0.028)
Fear or violence (is reason, 0/1)		0.029
		(0.028)
Racism (is reason, 0/1)		0.108***
		(0.031)
Having a different culture/speech (is reason, 0/1)		0.047
		(0.029)
Dislike of favela residents (is reason, 0/1)		0.029
		(0.029)
Nepotism (is reason, 0/1)		-0.061**
		(0.027)
Observations	2166	1496

Note: OLS estimates. Outcome is a dummy variable for whether the jobseeker expected at-or-above-median discrimination when predicting the audit study. See notes to Table A.1 for independent variable descriptions. Robust standard errors shown between parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A.4: Address Omission Experiment: Randomization Balance

	(1) Address Omission		(2) Known Address		(3) Status Quo		(1)-(2)		(1)-(3) Pairwise t-test		(2)-(3)	
Variable	N	Mean/(SE)	N	Mean/(SE)	N	Mean/(SE)	N	Mean difference	N	Mean difference	N	Mean difference
Jacarezinho resident (0/1)	434	0.191 (0.019)	437	0.185 (0.019)	432	0.176 (0.018)	871	0.006	866	0.015	869	0.009
Manguinhos resident (0/1)	434	0.028 (0.008)	437	0.025 (0.008)	432	0.028 (0.008)	871	0.002	866	-0.000	869	-0.003
Age	434	26.894 (0.301)	437	27.085 (0.303)	432	26.333 (0.299)	871	-0.191	866	0.561	869	0.751*
Male (0/1)	434	0.350 (0.023)	437	0.265 (0.021)	432	0.269 (0.021)	871	0.085***	866	0.082***	869	-0.003
White jobseeker (0/1)	434	0.247 (0.021)	437	0.206 (0.019)	432	0.234 (0.020)	871	0.041	866	0.013	869	-0.028
Some college (0/1)	434	0.071 (0.012)	437	0.050 (0.010)	432	0.069 (0.012)	871	0.021	866	0.002	869	-0.019
Completed regular high-school (0/1)	434	0.786 (0.020)	437	0.783 (0.020)	432	0.759 (0.021)	871	0.003	866	0.026	869	0.023
Working now (0/1)	434	0.327 (0.023)	437	0.320 (0.022)	432	0.331 (0.023)	871	0.007	866	-0.004	869	-0.011
Holds a formal job (0/1)	434	0.127 (0.016)	437	0.119 (0.016)	432	0.109 (0.015)	871	0.008	866	0.018	869	0.010
Ever worked (0/1)	434	0.744 (0.021)	437	0.730 (0.021)	432	0.692 (0.022)	871	0.014	866	0.052*	869	0.038
Actively searched last week (0/1)	434	0.546 (0.024)	437	0.533 (0.024)	432	0.514 (0.024)	871	0.013	866	0.032	869	0.019
Skill index	434	0.019 (0.028)	437	-0.078 (0.026)	432	-0.040 (0.027)	871	0.097**	866	0.058	869	-0.038
Expected discrimination predicting audit results	434	53.131 (1.773)	437	51.617 (2.655)	432	54.277 (1.412)	871	1.514	866	-1.145	869	-2.660
Reservation wage (USD)	433	256.716 (4.699)	437	252.654 (6.074)	432	249.924 (4.759)	870	4.062	865	6.793	869	2.731
F-test of joint significance (F-stat)								1.588*		1.391		0.913
F-test, number of observations								870		865		869

Note: Means, mean comparisons, and F-test for joint significance of differences in covariates across pairs of treatment arms. See notes to Table A.1 for variable descriptions. * p<0.1, ** p<0.05, *** p<0.01.

Table A.5: Information Experiment: Randomization Balance

	(1) Favela Info		(2) Full Info		(3) No Info		(1)-(2)		(1)-(3) Pairwise t-test		(2)-(3)	
Variable	N	Mean/(SE)	N	Mean/(SE)	N	Mean/(SE)	N	Mean difference	N	Mean difference	N	Mean difference
Jacarezinho resident (0/1)	247	0.219 (0.026)	227	0.198 (0.027)	216	0.157 (0.025)	474	0.020	463	0.061*	443	0.041
Manguinhos resident (0/1)	247	0.417 (0.031)	227	0.436 (0.033)	216	0.514 (0.034)	474	-0.019	463	-0.097**	443	-0.078
Age	247	27.773 (0.412)	227	26.463 (0.394)	216	26.773 (0.414)	474	1.311**	463	1.000*	443	-0.311
Male (0/1)	247	0.308 (0.029)	227	0.313 (0.031)	216	0.287 (0.031)	474	-0.005	463	0.021	443	0.026
White jobseeker (0/1)	247	0.206 (0.026)	227	0.220 (0.028)	216	0.204 (0.027)	474	-0.014	463	0.003	443	0.017
Some college (0/1)	247	0.089 (0.018)	227	0.070 (0.017)	216	0.079 (0.018)	474	0.019	463	0.010	443	-0.008
Completed regular high-school (0/1)	247	0.846 (0.023)	227	0.775 (0.028)	216	0.847 (0.025)	474	0.071*	463	-0.001	443	-0.072*
Working now (0/1)	247	0.304 (0.029)	227	0.260 (0.029)	216	0.287 (0.031)	474	0.044	463	0.017	443	-0.027
Holds a formal job (0/1)	247	0.162 (0.023)	227	0.115 (0.021)	216	0.125 (0.023)	474	0.047	463	0.037	443	-0.010
Ever worked (0/1)	247	0.826 (0.024)	227	0.758 (0.029)	216	0.769 (0.029)	474	0.068*	463	0.057	443	-0.011
Actively searched last week (0/1)	247	0.449 (0.032)	227	0.427 (0.033)	216	0.394 (0.033)	474	0.022	463	0.056	443	0.034
Skill index	247	0.040 (0.037)	227	-0.044 (0.038)	216	0.048 (0.039)	474	0.085	463	-0.007	443	-0.092*
Expected discrimination predicting audit results	247	33.297 (4.568)	227	28.961 (4.864)	216	37.886 (2.372)	474	4.336	463	-4.589	443	-8.924*
Reservation wage (USD)	247	249.212 (4.708)	227	244.605 (6.522)	216	244.344 (5.455)	474	4.607	463	4.868	443	0.261
F-test of joint significance (F-stat)								0.894		1.154		0.998
F-test, number of observations								474		463		443

Note: Means, mean comparisons, and F-test for joint significance of differences in covariates across pairs of treatment arms. See notes to Table A.1 for variable descriptions. * p<0.1, ** p<0.05, *** p<0.01.

Table A.6: Interview Experiment: Randomization Balance

Variable	(1) Name-Only Mean/(SE)	(2) Name-and-Address Mean/(SE)	(1)-(2) Pairwise t-test Mean difference
Address Omisson	0.227	0.243	-0.015
	(0.028)	(0.030)	
Known Address	0.218	0.233	-0.014
	(0.028)	(0.030)	
Favela Info	0.141	0.153	-0.013
	(0.024)	(0.025)	
Full Info	0.150	0.134	0.016
	(0.024)	(0.024)	
Jacarezinho resident (0/1)	0.232	0.173	0.059
	(0.029)	(0.027)	
Manguinhos resident (0/1)	0.200	0.163	0.037
	(0.027)	(0.026)	
Age	25.768	25.847	-0.078
	(0.409)	(0.390)	
Male (0/1)	0.268	0.262	0.006
	(0.030)	(0.031)	
White jobseeker (0/1)	0.236	0.238	-0.001
	(0.029)	(0.030)	
Some college (0/1)	0.064	0.079	-0.016
	(0.016)	(0.019)	
Completed regular high-school (0/1)	0.786	0.767	0.019
	(0.028)	(0.030)	
Working now (0/1)	0.164	0.104	0.060*
	(0.025)	(0.022)	
Holds a formal job (0/1)	0.050	0.045	0.005
	(0.015)	(0.015)	
Ever worked (0/1)	0.764	0.708	0.056
	(0.029)	(0.032)	
Actively searched last week (0/1)	0.627	0.673	-0.046
	(0.033)	(0.033)	
Skill index	0.032	0.009	0.023
	(0.037)	(0.040)	
Expected discrimination predicting audit results	44.508	45.582	-1.074
	(4.537)	(3.169)	
Reservation wage (USD)	232.233	231.667	0.565
	(3.830)	(3.918)	
F-test of joint significance (F-stat)			0.814
Number of observations	220	202	422

Note: Means, mean comparisons, and F-test for joint significance of differences in covariates across pairs of treatment arms. See notes to Table A.1 for variable descriptions. * p<0.1, ** p<0.05, *** p<0.01.

Table A.7: Comparison of Samples Across the Three Experiments

	(1) Address Omission Experiment		(2) Information Experiment		(3) Interview Experiment		(1)-(2)		(1)-(3) Pairwise t-test		(2)-(3)	
Variable	N	Mean/(SE)	N	Mean/(SE)	N	Mean/(SE)	N	Mean difference	N	Mean difference	N	Mean difference
Jacarezinho resident (0/1)	1303	0.184	690	0.193	422	0.204	1993	-0.009	1725	-0.020	1112	-0.011
		(0.011)		(0.015)		(0.020)						
Manguinhos resident (0/1)	1303	0.027	690	0.454	422	0.182	1993	-0.427***	1725	-0.156***	1112	0.271***
		(0.004)		(0.019)		(0.019)						
Age	1303	26.772	690	27.029	422	25.806	1993	-0.257	1725	0.966***	1112	1.223***
		(0.174)		(0.236)		(0.283)						
Male (0/1)	1303	0.295	690	0.303	422	0.265	1993	-0.008	1725	0.029	1112	0.037
		(0.013)		(0.018)		(0.022)						
White jobseeker (0/1)	1303	0.229	690	0.210	422	0.237	1993	0.019	1725	-0.008	1112	-0.027
		(0.012)		(0.016)		(0.021)						
Working now (0/1)	1303	0.326	690	0.284	422	0.135	1993	0.042*	1725	0.191***	1112	0.149***
		(0.013)		(0.017)		(0.017)						
Holds a formal job (0/1)	1303	0.118	690	0.135	422	0.047	1993	-0.017	1725	0.071***	1112	0.087***
		(0.009)		(0.013)		(0.010)						
Ever worked (0/1)	1303	0.722	690	0.786	422	0.737	1993	-0.063***	1725	-0.015	1112	0.049*
		(0.012)		(0.016)		(0.021)						
Actively searched in the last 7 days	1303	0.531	690	0.425	422	0.649	1993	0.106***	1725	-0.118***	1112	-0.225***
		(0.014)		(0.019)		(0.023)						
High skill (0/1)	1303	0.487	690	0.558	422	0.559	1993	-0.071***	1725	-0.073***	1112	-0.001
		(0.014)		(0.019)		(0.024)						
Expected discrimination predicting audit results	1303	53.003	690	33.307	422	45.022	1993	19.696***	1725	7.981***	1112	-11.715***
		(1.166)		(2.406)		(2.807)						
Reservation wage (USD)	1303	253.106	690	246.173	422	231.962	1993	6.934	1725	21.144***	1112	14.211***
		(3.011)		(3.215)		(2.736)						
F-test of joint significance (F-stat)								80.830***		22.027***		16.382***
F-test, number of observations								1993		1725		1112

Notes: Pair-wise comparisons of average baseline characteristics across experiments. See notes to Table A.1 for details on variables. * p<0.1, ** p<0.05, *** p<0.01.

Table A.8: Pearson Correlation Coefficients Between Interview Performance Indicators

	Overall (Interv.)	Calm (Interv.)	Professional (Interv.)	Overall (Self)	Calmn (Self)	Professional (Self)
Overall (Interv.)	1					
Calm (Interv.)	0.48 (0.04)	1				
Professional (Interv.)	0.71 (0.03)	0.29 (0.05)	1			
Overall (Self)	0.29 (0.05)	0.27 (0.05)	0.15 (0.05)	1		
Calmn (Self)	0.21 (0.05)	0.27 (0.05)	0.10 (0.05)	0.24 (0.05)	1	
Professional (Self)	0.11 (0.05)	0.16 (0.05)	0.06 (0.05)	0.42 (0.04)	-0.01 (0.05)	1

Note: Variables are the components of the interview performace indexes, and "Interv." refers to interviewer assessments. Standard errors between parenthesis.

Table A.9: Interview Treatment Effects by Treatment Conditions Before Interview

	(1) Aggregate performance index	(2) Aggregate performance index	(3) Interviewer-assessed performance index	(4) Interviewer-assessed performance index	(5) Self-assessed performance index	(6) Self-assessed performance index
Name-Only× Status Quo × No Info (N=129)	0.23**	0.12	0.36***	0.25*	0.10	-0.01
	(0.10)	(0.10)	(0.14)	(0.13)	(0.12)	(0.12)
Name-Only × Status Quo × Favela Info (N=55)	0.26	0.30*	0.36	0.37*	0.16	0.22
	(0.18)	(0.17)	(0.23)	(0.21)	(0.20)	(0.20)
Name-Only × Status Quo × Full Info (N=44)	0.07	0.07	-0.17	-0.13	0.30*	0.27
	(0.16)	(0.16)	(0.23)	(0.23)	(0.18)	(0.18)
Name-Only × Address Omission × No Info (N=90)	-0.09	-0.07	-0.30	-0.27	0.13	0.12
	(0.13)	(0.12)	(0.20)	(0.19)	(0.13)	(0.12)
Name-Only × Known Address × No Info (N=81)	0.08	0.00	0.06	-0.05	0.11	0.06
	(0.13)	(0.14)	(0.17)	(0.18)	(0.16)	(0.16)
Name-Only × Other Combinations (N=23)	0.25	0.34	-0.12	-0.03	0.61**	0.71**
	(0.25)	(0.25)	(0.40)	(0.41)	(0.27)	(0.30)
Observations	422	422	422	422	422	422
Controls	No	Yes	No	Yes	No	Yes
P-value for no effect on more externally valid subsample	0.03	0.11	0.01	0.04	0.51	0.53
P-value for no effect on less externally valid subsample	0.80	0.91	0.40	0.49	0.23	0.34
P-value for equal effect on more vs. less externally valid	0.08	0.19	0.00	0.01	0.53	0.33

Note: OLS estimates for the effects of *Name-Only* on interview performance indexes for subgroups defined by pre-interview-stage treatments. There is no coefficient for *Name-Only* as all jobseekers were assigned to some application-stage treatment. Regressions fully control for treatment assignment in previous experiments, and also include the variables on balance tables as control in even columns. On the bottom rows, "more externally valid subsample" refers to the first two coefficients in each line, and "less externally valid subsample" refers to all other coefficients (all weighted by share of total participants under the application-stage assignment). Robust standard errors between parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A.10: No Effect Heterogeneity by Race in the Information Experiment Endline

	(1) Responded to endline (0/1)	(2) Exp. discrimination (categorical, 1-4)	(3) # of sent apps (categorical, 1-4)
Favela Info	0.01	0.12	-0.08
	(0.05)	(0.11)	(0.15)
Full Info	0.02	-0.15	-0.01
	(0.05)	(0.11)	(0.16)
Level Info × white	0.05	-0.27	0.30
	(0.12)	(0.25)	(0.32)
Full Info × white	-0.03	0.14	0.14
	(0.12)	(0.28)	(0.35)
White jobseeker (0/1)	-0.01	0.10	-0.09
	(0.09)	(0.20)	(0.24)
Observations	690	389	389
No Info Mean among non-white jobseekers	0.55	2.27	2.55
Favela=Full among white jobseekers *p*	0.50	0.50	0.77

Note: See notes to Table 1.

Table A.11: Table 3 Including Interactions

	Experiment-specific outcomes			
	(1)	(2)	(3)	(4)
Agg. interv. performance (SD)	0.193*	0.218***	0.111	0.161*
Name-Only $\times h_i$	(0.101)	(0.078)	(0.088)	(0.085)
Show-up (%)	7.484	−0.189	−5.633	5.414
Address Omission $\times h_i$	(4.918)	(3.981)	(3.874)	(4.018)
Show-up (%)	−4.640	0.209	0.854	0.926
Known Address $\times h_i$	(4.166)	(3.968)	(4.370)	(3.810)
Show-up (%)	−13.432**	−0.777	−10.084**	11.080**
Full Info $\times h_i$	(5.817)	(5.001)	(4.778)	(5.307)
Show-up (%)	3.437	2.684	−7.762	11.413**
Favela Info $\times h_i$	(6.648)	(5.358)	(5.694)	(5.462)
Heterogeneity variable h	White	High $\mathbb{E}$[disc]	Male	High Skill
Any heterogeneity by h, p-value	0.019	0.129	0.065	0.012
Heterogeneity by h in $\mathbb{E}$[address visib.] treatments, p-value	0.047	0.039	0.299	0.094
Heterogeneity by h in information treatment, p-value	0.051	0.865	0.036	0.01
Clusters	2032	2032	2032	2032
Observations	2,414	2,414	2,414	2,414

Note: This table reports the same stacked OLS estimates than Table 3, but the model also includes interactions with the variables included in the balance checks. Note that this implies all columns are actually the same regression.

Table A.12: Two-stage Least Squares Estimates of The Effect of the Expected Callback Rate and Antifavela 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 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 for each audit study neighborhood, 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.

Table A.13: Race Correlates: College Attendance, Work Experience, Age, and Skill

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Jacarezinho resident (0/1)	-0.08***	-0.06***											
	(0.02)	(0.02)											
Manguinhos resident (0/1)	-0.05**		-0.02										
	(0.02)		(0.02)										
Age	-0.00***			-0.00***									
	(0.00)			(0.00)									
Male (0/1)	0.03				0.03								
	(0.02)				(0.02)								
Working now (0/1)	-0.06**					-0.05***							
	(0.02)					(0.02)							
Holds a formal job (0/1)	0.03						-0.02						
	(0.03)						(0.03)						
Ever worked (0/1)	-0.02							-0.05**					
	(0.02)							(0.02)					
Actively searched in the last 7 days	-0.01								-0.02				
	(0.02)								(0.02)				
High skill (0/1)	0.08***									0.07***			
	(0.02)									(0.02)			
Prior Disc (%)	-0.00										-0.00		
	(0.00)										(0.00)		
Reservation wage (USD)	0.00											0.00	
	(0.00)											(0.00)	
Expects high discrimination (0/1)	-0.01												-0.03
	(0.02)												(0.02)
Observations	2167	2167	2167	2167	2167	2167	2167	2167	2167	2167	2167	2167	2167

Note: * p<0.1, ** p<0.05, *** p<0.01. Outcome is a dummy (0/1) for the jobseekers being white. Robust standard errors shown between parenthesis.

Table A.14: Effects on Interview Performance for the High Expected Discrimination and White Individuals Are Robust to Including Other Interacted Covariates

	(1) Aggregated performance index	(2) Aggregated performance index	(3) Aggregated performance index	(4) Interviewer-assessed performance	(5) Interviewer-assessed performance	(6) Interviewer-assessed performance	(7) Self-assessed performance	(8) Self-assessed performance	(9) Self-assessed performance
Name-Only	-0.02	-0.07	0.07	-0.11	-0.18	-0.04	0.07	0.03	0.18
	(0.09)	(0.09)	(0.38)	(0.13)	(0.14)	(0.52)	(0.09)	(0.10)	(0.46)
Name-Only × High Exp. Disc.	0.24**	0.24**	0.28**	0.33**	0.32*	0.34**	0.16	0.16	0.22
	(0.12)	(0.12)	(0.12)	(0.17)	(0.16)	(0.17)	(0.13)	(0.13)	(0.14)
Name-Only × White jobsee.		0.24*	0.31**		0.29	0.35*		0.19	0.27*
		(0.14)	(0.14)		(0.20)	(0.19)		(0.15)	(0.16)
Observations	422	422	422	422	422	422	422	422	422
Other Interactions?	No	No	Yes	No	No	Yes	No	No	Yes

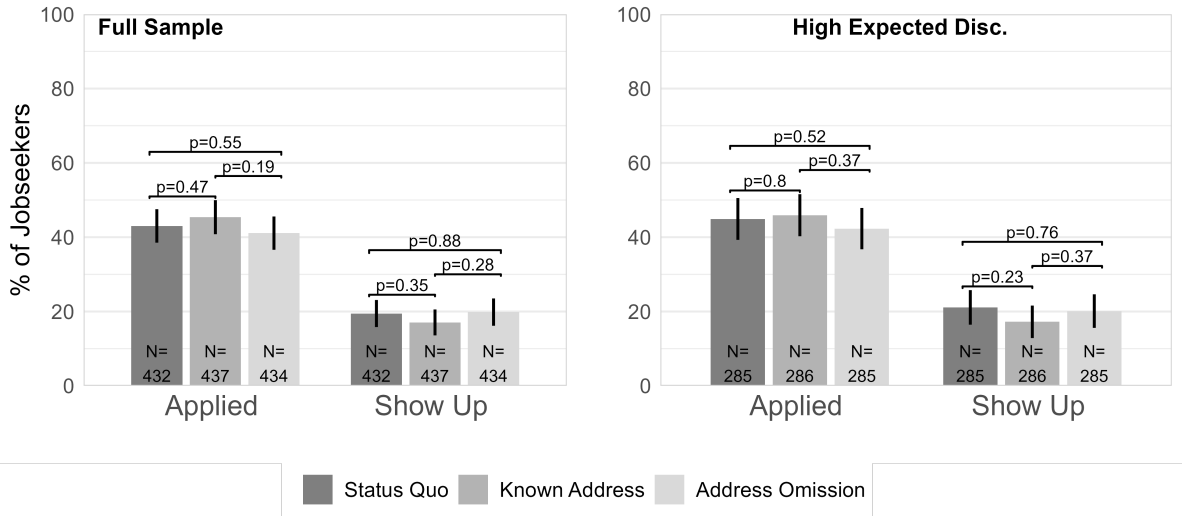
Notes: Table displays OLS coefficients for regressions of the form $y_i = \alpha + \beta Name\text{-}Only + \gamma Name\text{-}OnlyX_i + \mu X_i + \varepsilon$, where X_i includes a high expected discrimination dummy (columns 1, 4, and 7), that same dummy plus a dummy for race (columns 2, 5, and 8), or all variables in Table A.6 (columns 3, 6, and 9). * p<0.1, ** p<0.05, *** p<0.01.

Table A.15: Effects on Interview Performance Concentrate on White Jobseekers with Names that Are Not Distictive of Low SES

	(1) Aggregated performance index	(2) Aggregated performance index	(3) Interviewer-assessed performance	(4) Interviewer-assessed performance	(5) Self-assessed performance	(6) Self-assessed performance
Name-Only× nonlow SES name × white	0.41***	0.37***	0.41**	0.36**	0.41***	0.37***
	(0.13)	(0.12)	(0.19)	(0.17)	(0.14)	(0.14)
Name-Only× nonlow SES name × nonwhite	0.07	0.05	0.03	0.01	0.11	0.10
	(0.08)	(0.08)	(0.11)	(0.11)	(0.09)	(0.09)
Name-Only× low SES name × white	-0.01	0.03	-0.03	-0.02	-0.00	0.07
	(0.33)	(0.28)	(0.44)	(0.39)	(0.32)	(0.28)
Name-Only× low SES name × nonwhite	0.06	-0.00	-0.02	-0.09	0.14	0.08
	(0.11)	(0.11)	(0.16)	(0.15)	(0.13)	(0.13)
Observations	422	422	422	422	422	422
Controls	No	Yes	No	Yes	No	Yes

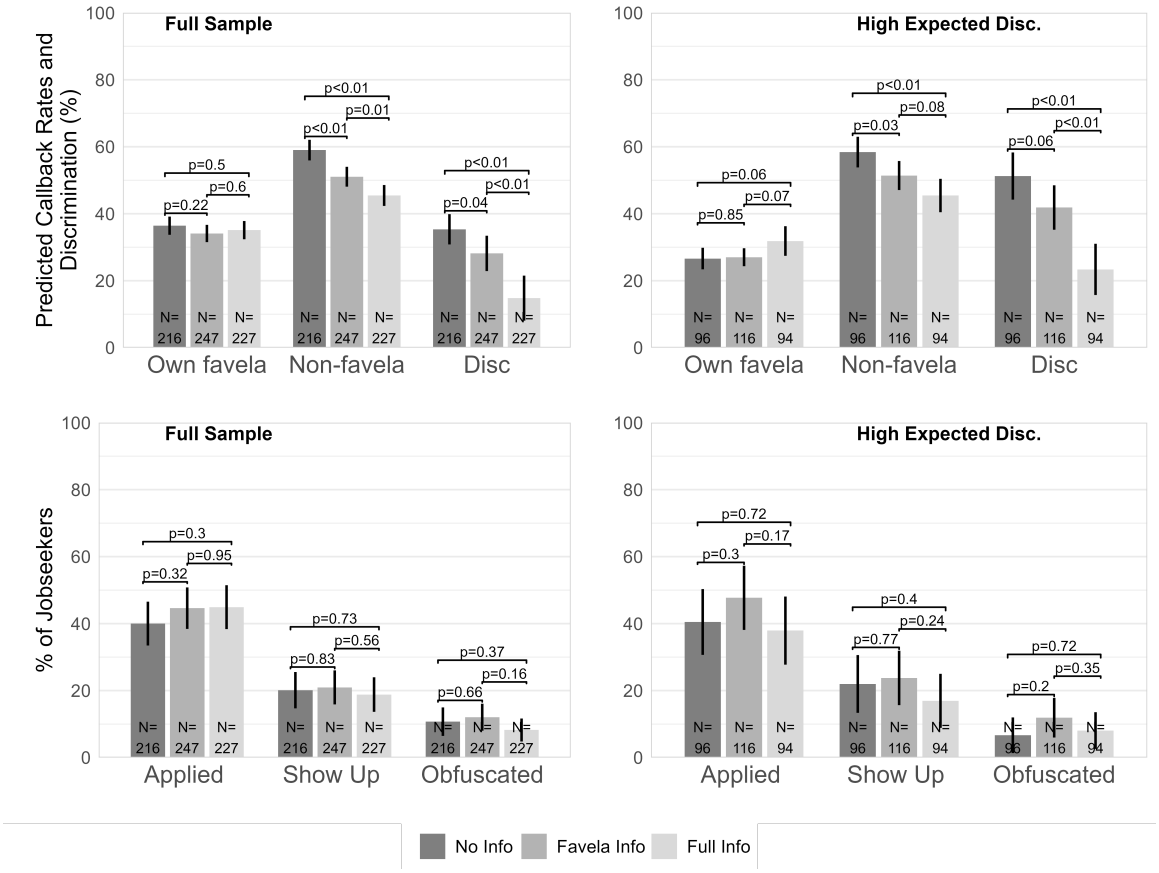
Notes: Table displays OLS estimates of treatment effects on interview performance by each combination of race and type of first name, without an omitted category (so *Name-Only* by itself is not included). Controls by race and type of name are always included, and all variables in Table A.6 are also included as controls in even columns. Names were classified into those that might or might not be distinctive of low SES by first asking ChatGPT 4o and then manually revising the classification, in light of our knowledge of the context. * p<0.1, ** p<0.05, *** p<0.01.

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



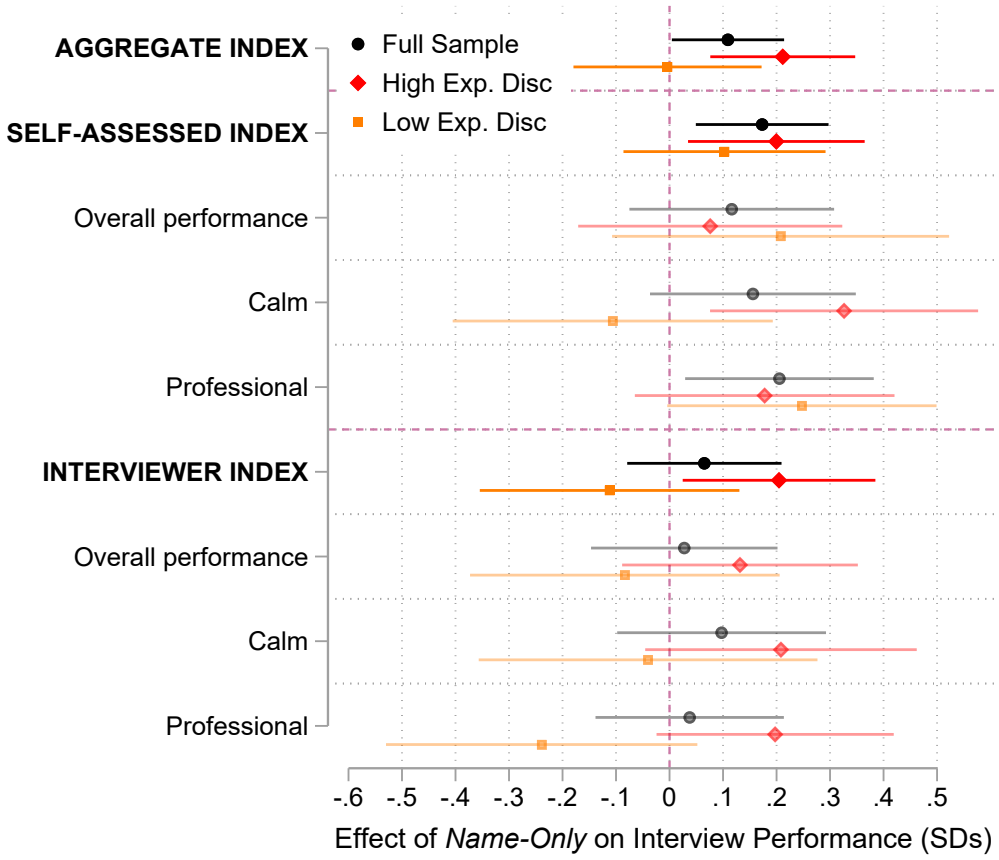
Notes: The double-lasso procedure could pick controls including expected discrimination, race, gender, residing in Maré, education, communication skills, math test score, having ever worked, working now, age, and whether the jobseeker has looked for a job in the last 7 days. After picking controls, we demean them and interacting them with treatment status, so the bar heights reflect control-corrected means and differences reflect ATE between arms. Variance-covariance matrix was estimated using the HC1 approach since the extra controls and sample-splitting generate ill-defined entries; see notes for Figure 4 for other details.

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



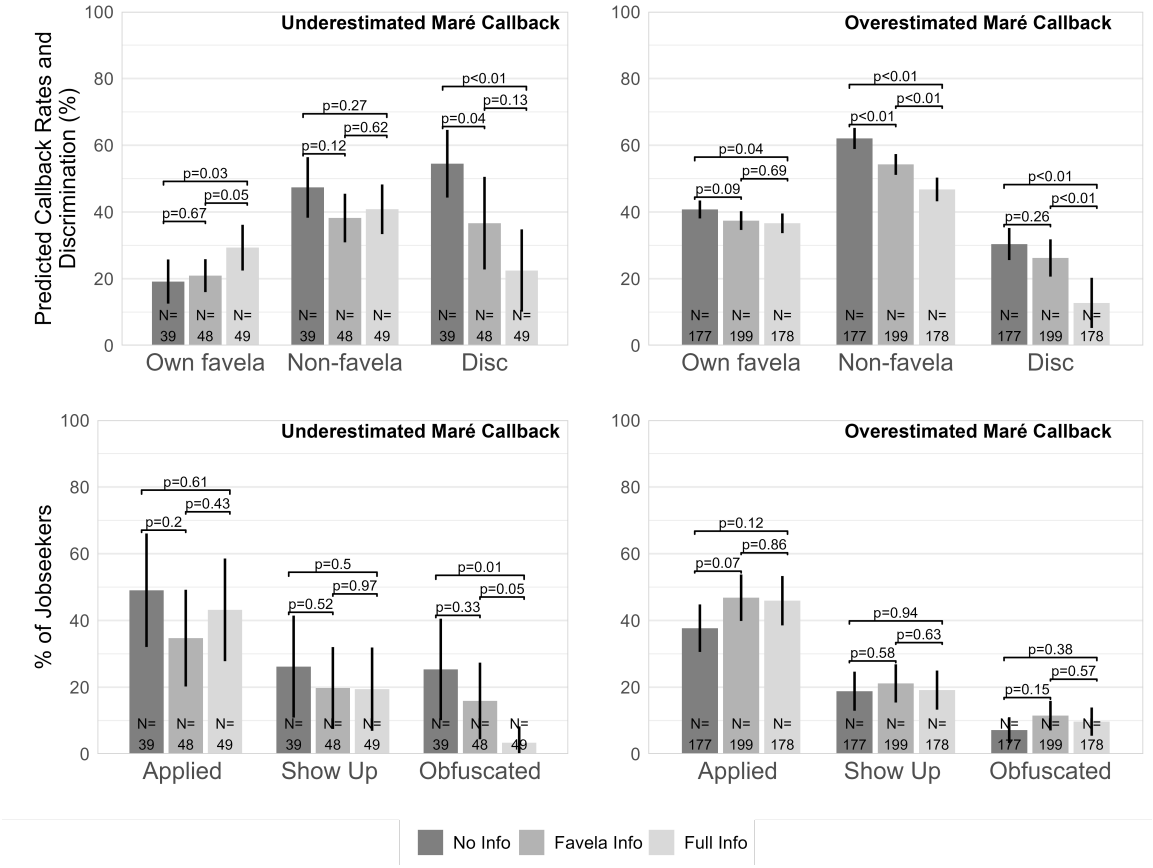
Note: The double-lasso procedure could pick controls among the variables included in balance checks. After picking controls, we demean them and interacting them with treatment status, so the bar heights reflect control-corrected means and differences reflect ATE between arms. Variance-covariance matrix was estimated using the HC1 approach since the extra controls and sample-splitting generate ill-defined entries; see notes for Figure 5 for other details.

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



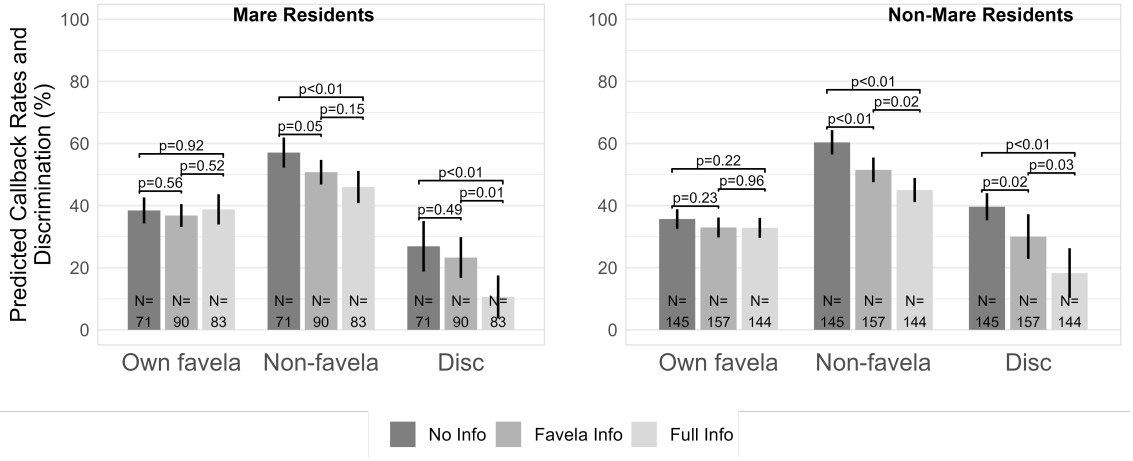
Note: The graph shows treatment effect estimates (using double-lasso selected controls) for the full sample and conditional on expecting below or at-or-above median discrimination at baseline. The double-lasso procedure could pick controls including expected discrimination, race, gender, residing in Maré, education, communication skills, math test score, computer skills, having ever worked, working now, age, and whether the jobseeker has looked for a job in the last 7 days. The interview performance outcomes are listed on the left-hand side and described in Section 3.2.3. 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 Jobseekers Initially Under- or Overestimated the Favela Callback Rate



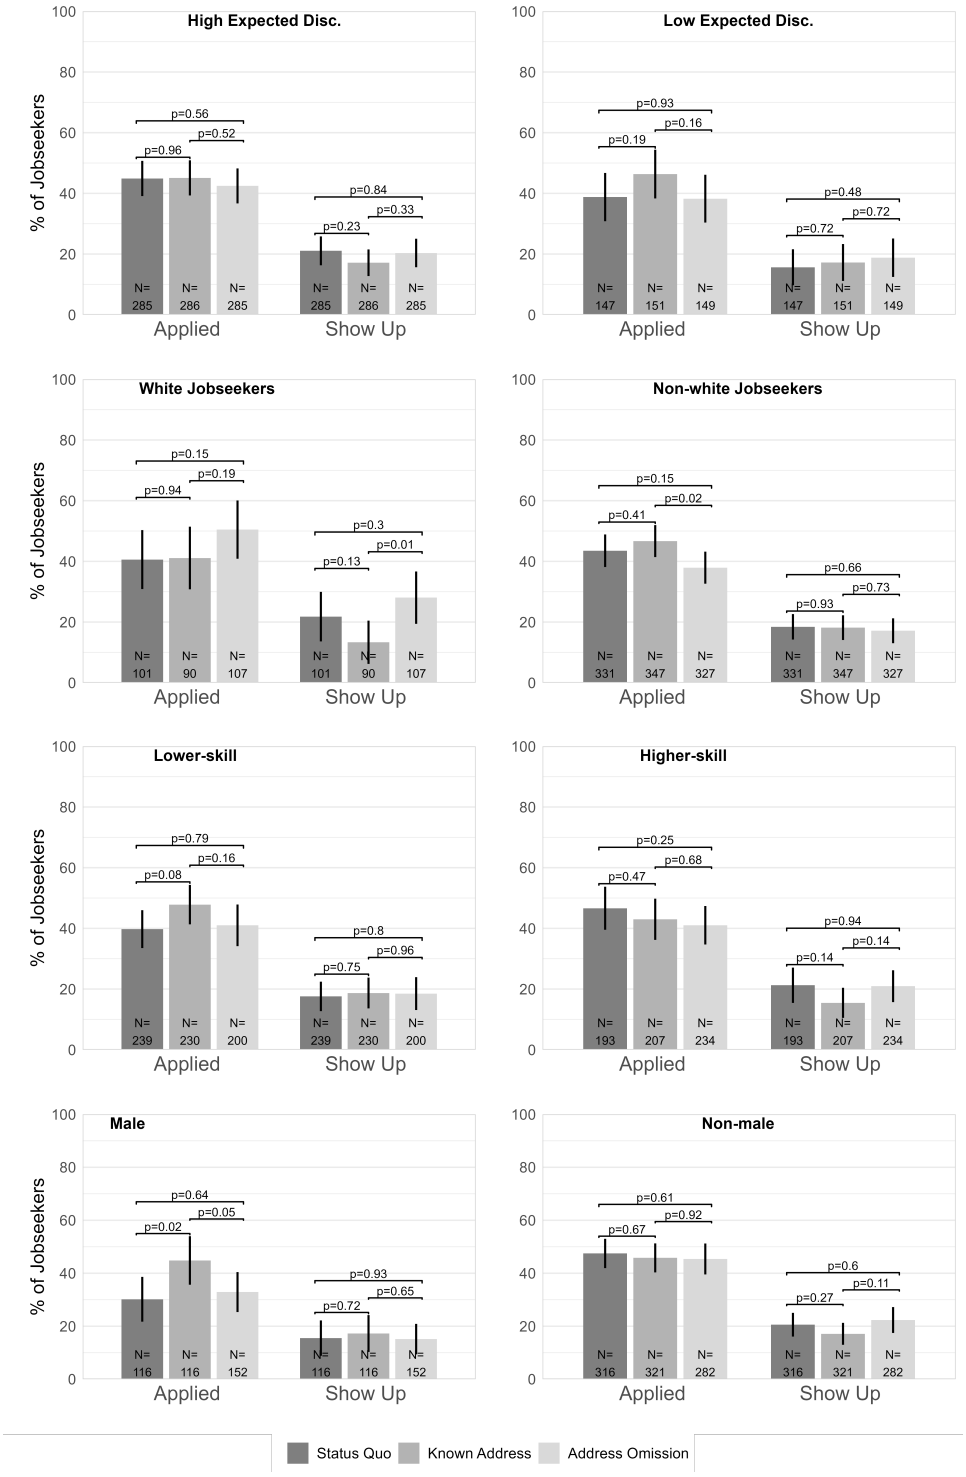
Note: Same as in Figure 5, but splitting the same on those who under- or over-estimated the callback rate for the favela neighborhood in the audit study. This makes it easier to see that how jobseekers adjust beliefs about callback rates for their own favela according to the direction of the information received. Note also that *Full Info* decreases the obfuscation rate for those initially too pessimistic about favela callback rates, consistent with obfuscation becoming pointless once observing there are little returns to it at the callback stage.

Figure A.5: Belief Update in Information Experiment Occurs for Maré and Non-Maré Residents



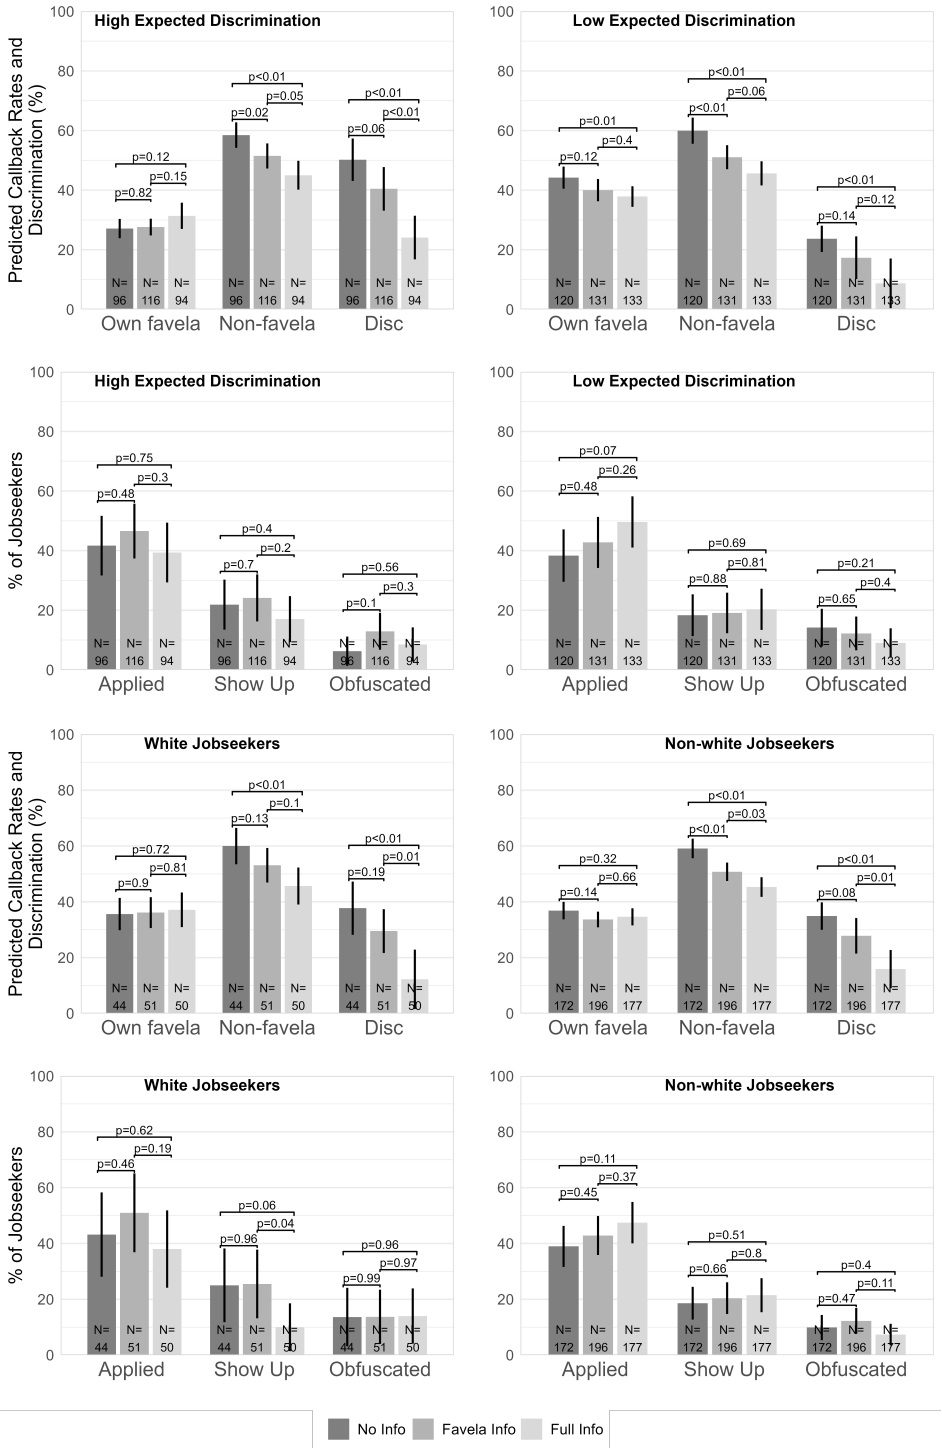
Notes: The callback rates revealed in the information experiment were those found in the audit study, for Maré and Bonsucesso. Residents of Manguinhos and Jacarezinho make similar belief updates as the Maré residents, suggesting that they extrapolate similarly from the audit findings. See notes to Figure 5 for details on outcomes and figure features.

Figure A.6: Conditional Treatment Effects – Address Omission Experiments



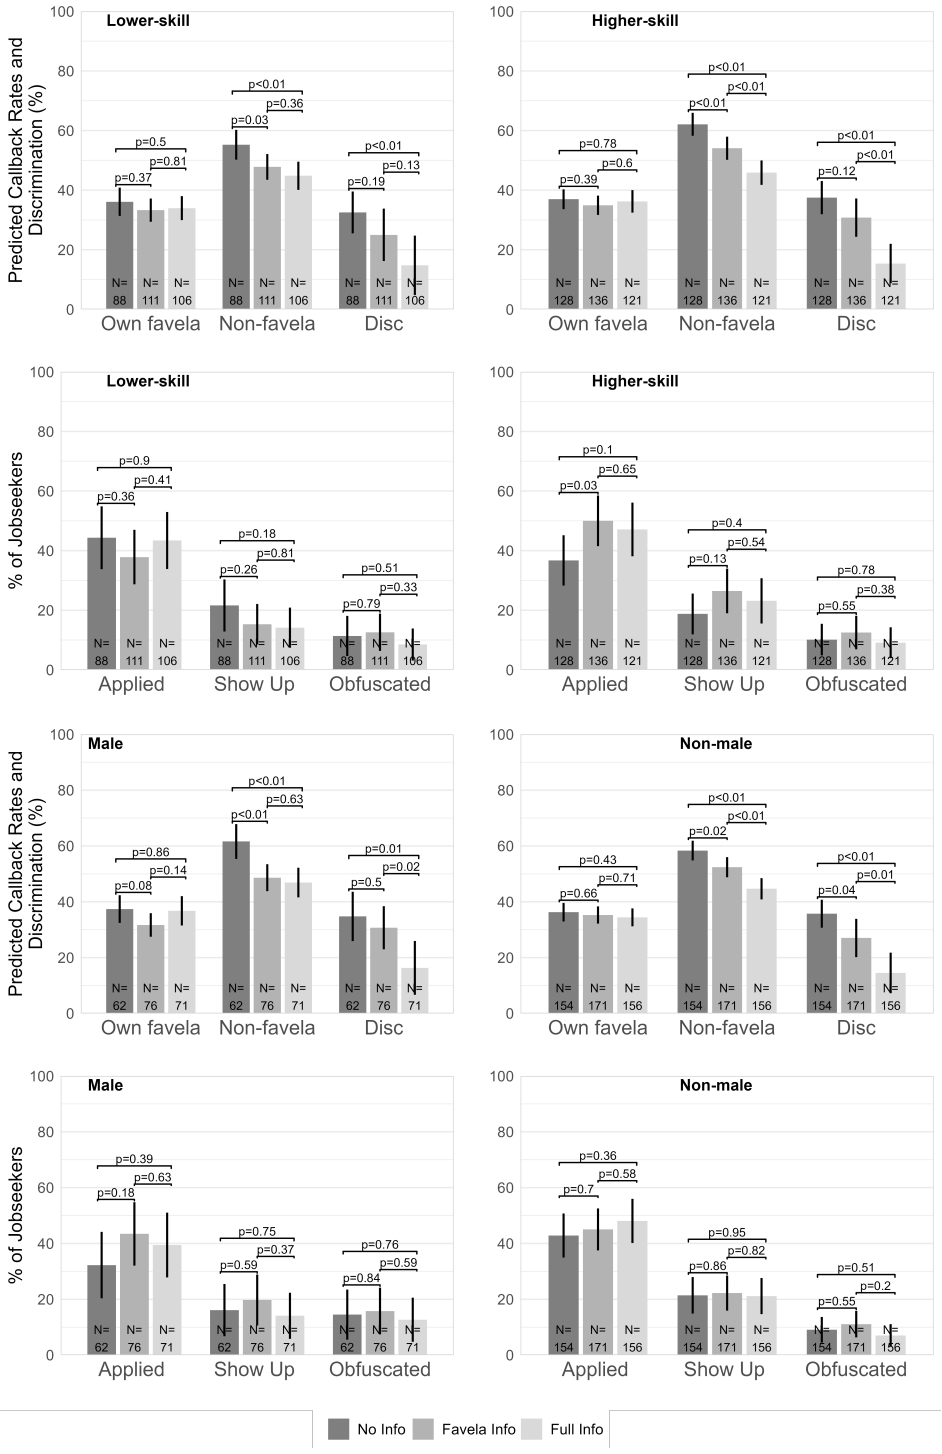
Notes: Figures show job application and interview attendance rates conditional on binary sample splits (by the four preregistered heterogeneity break-downs). See notes to Figure 4 for details and Section 4.4 for a systematic discussion of treatment effect heterogeneity.

Figure A.7: Conditional Treatment Effects – Information Experiment, Part 1



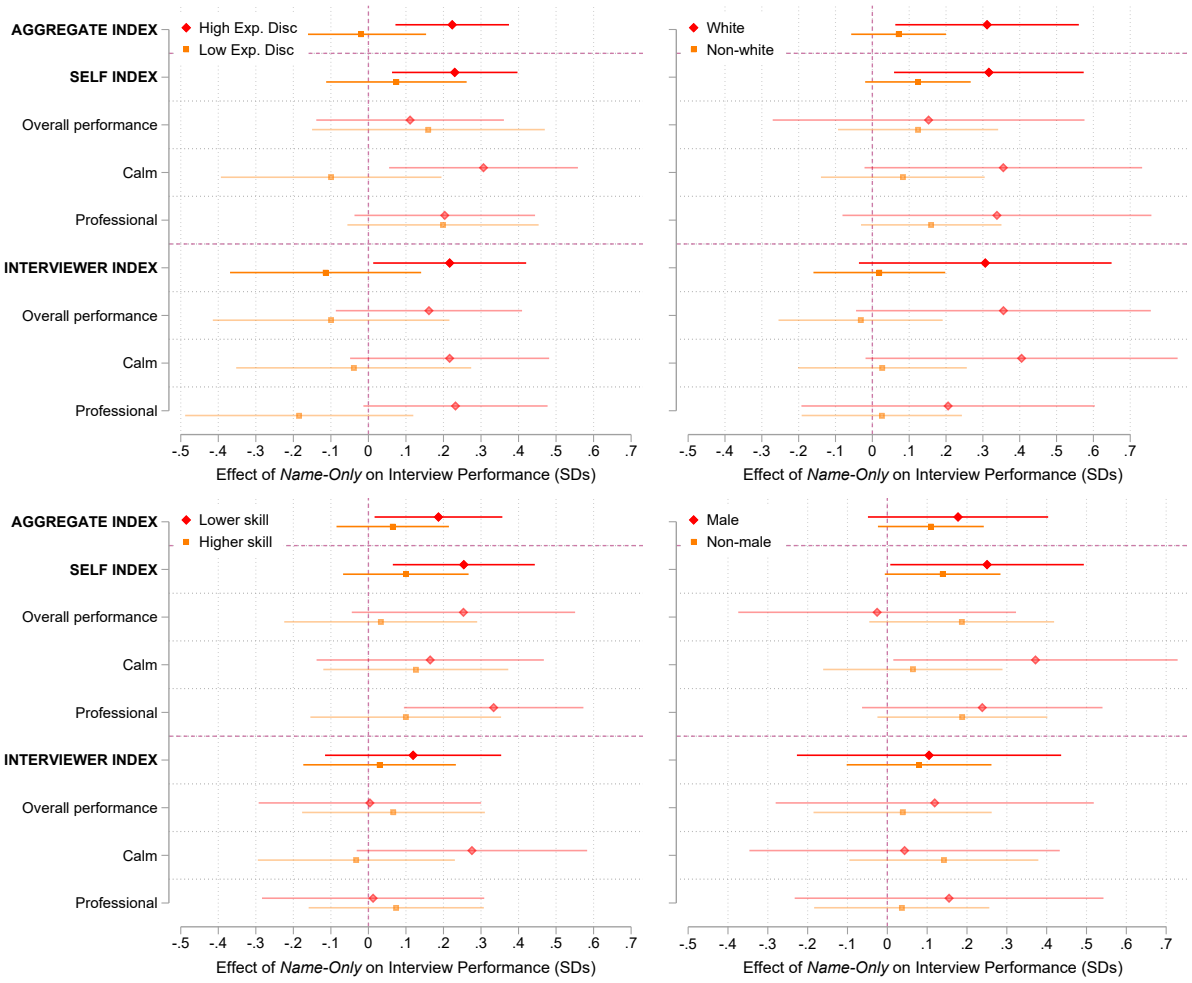
Notes: Figures show job application and interview attendance rates conditional on binary sample splits (by the four preregistered heterogeneity break-downs). See notes to Figure 5 for details and Section 4.4 for a systematic discussion of treatment effect heterogeneity.

Figure A.8: Conditional Treatment Effects – Information Experiment, Part 2



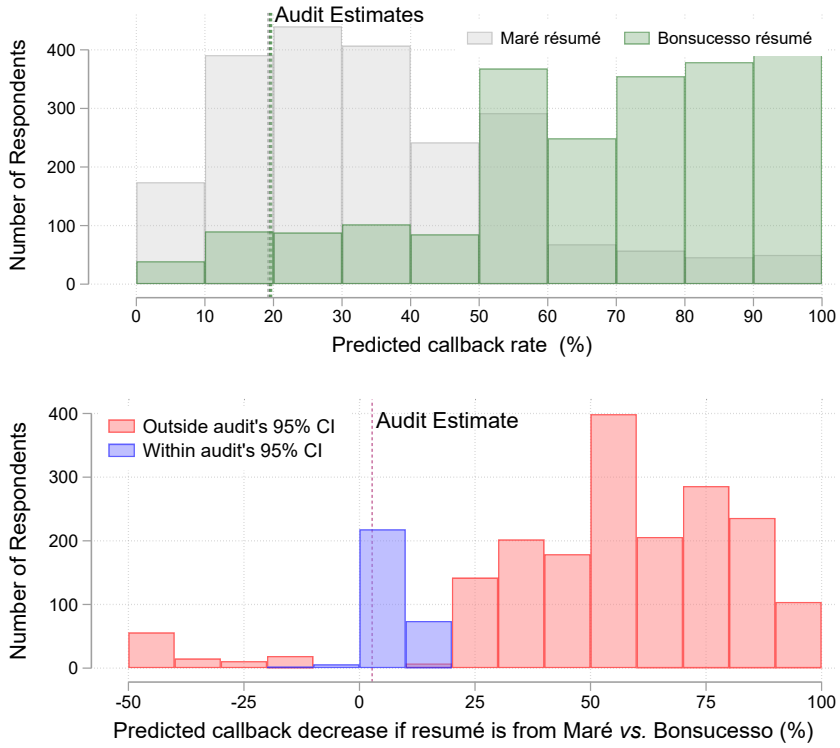
Notes: Figures show job application and interview attendance rates conditional on binary sample splits (by the four preregistered heterogeneity break-downs). See notes to Figure 5 for details and Section 4.4 for a systematic discussion of treatment effect heterogeneity.

Figure A.9: Conditional Treatment Effects of *Name-Only*



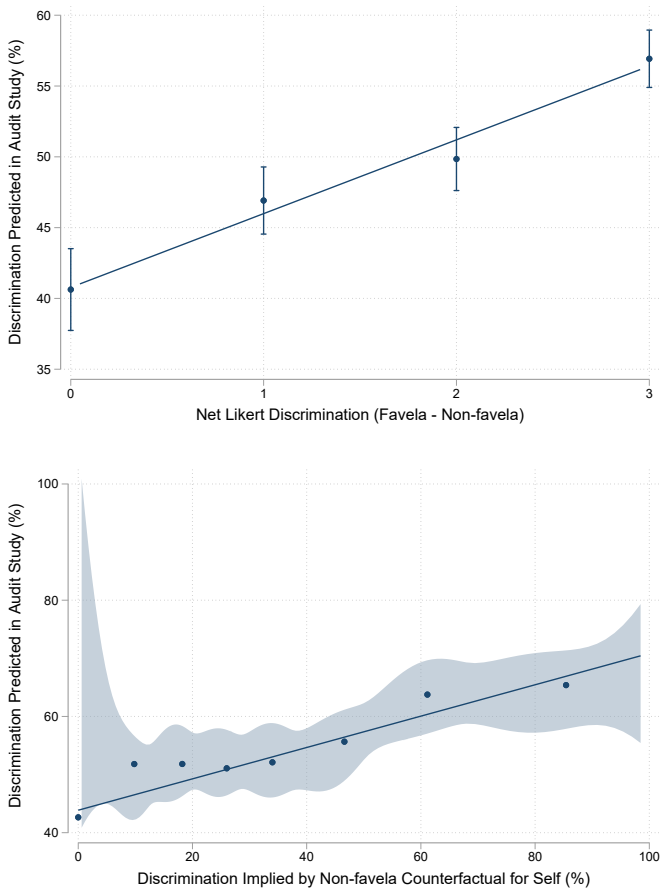
Notes: Preregistered heterogeneity break-downs for the interview experiment. See notes to Figure 6 for details and Section 4.4 for a systematic discussion of treatment effect heterogeneity.

Figure A.10: Predicted *vs.* Actual Discrimination Rates Using Only Beliefs About the Audit Study Neighborhoods



Notes: Same as in Figure 2, but using declared beliefs regarding Maré and Bonsucesso instead of one's own favela and its adjacent nonfavela.

Figure A.11: Predicted Audit Study Discrimination Correlates with Other Proxies 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 Likert-scale answers of how much employers discriminate against individuals in one's favela and adjacent nonfavela (from no discrimination to a lot), converting them into integers from one to four, and then taking the difference between neighborhoods. We calculate the discrimination for the counterfactual self by comparing the beliefs about one's job-finding probability over the next six months to predictions about "somebody just like you, but from [adjacent nonfavela neighborhood]".

B Deviations from the Pre-Analysis Plan

Our preregistration can be accessed at https://www.socialscienceregistry.org/trials/11041. On that page, we also discuss our analysis plan with respect to effect heterogneneity. Below, we list our deviations from that plan.

- As mentioned in footnote 9, the one major deviation from the initial preregistration was the introduction of the information experiment. We amended the preregistration, explaining our reasoning, before inviting the N=690 participants in that experiment to apply.
- The first heterogeneity analysis mentioned in our preregistration is by expected discrimination level. In that same bullet, we also mention considering heterogeneity analyses by attitudes such as "how bothered one is by discrimination" and "whether the possibility of being discriminated against in the hiring process is motivating or discouraging." As we dropped these questions from our survey when introducing the information experiment, we also drop this more exploratory analysis.
- We preregistered 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 (one dummy for completing regular high school and one dummy for having some college-level education) and a measure of communication skills (Likert-scale, coded by the surveyor after the end of each survey) to build our skill index.
- We list effort in applications, measured by length and quality of the optional cover letter applicants could write in the online application form, as primary outcomes. For brevity, we omit them from the main text and present them here.
- 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 scheduled for the day was too low for make up a single strata.
- We initially planned to stratify the randomization in the interview experiment by predicted discrimination level *and* previous treatment assignments. During the implementation, we only stratified by the discrimination level. That is because, given lower-than-expected interview show-up rates, the batch sizes had to be smaller to keep a constant flow of participants into interviews.

- The preregistration mentions an endline phone survey with participants of the Address Omission Experiment. Its main purpose was to quality-check data collected by surveyors. As the experiment progressed, we added to this survey questions on whether the jobseeker applied to other jobs besides the ones in this study (see Table B.1). Later, we also added questions to verify what information job seekers believed that the HR firm in this study had about them prior to the application invite. We only asked 370 jobseekers about the latter before shifting to the Information Experiment. We do not find evidence of a first stage on expected address visibility, but we believe that is due to noise and imperfect recall: These phone calls took place at least four weeks after the invite to apply, and only about 60% believe (when asked a placebo question) that the HR firm knew candidates' phone numbers – showing that recall is quite imperfect (Table B.1).
- 184 out of the main sample of 2,167 jobseekers were randomly assigned both to a condition in the Address Omission Experiment and to a condition in the Information Experiment. Out of those 184, 151 participated in a "pilot" version of the Information Experiment, in which we did not add the callbacks that could not be matched to a specific job ad (so information was that the callback rates were around 16% for both neighborhoods). The remaining 33 received the same information described in the main text. 10 of those 33 were randomly assigned to the *Status Quo* condition, so those 10 are part of the information experiment in the main text. To summarize, the **2,167** sample is made of: **1,303** who received only Address Omission assignments, **680** who received only Information assignments, and **184** who received both (10 of which were prerandomization identical to the sample who only received Information assignments, so we pool them). See Appendix Section B.2

B.1 Treatment Effects on Secondary Outcomes

Table B.1: Treatment Effects on Secondary Outcomes in the Address Omission Experiment

	(1) Clicked application link (%)	(2) Words in cover letter	(3) Cover letter quality (0/100)	(4) Years of experience declared	(5) Declared experience in favela (0/1)	(6) Participated in endline (0/1)	(7) Applied for another job (0/1)	(8) Thought HR knew address before applying (0/1)	(9) Thought HR knew phone before applying (0/1)
Address Omission	-2.13	5.13	0.65	0.15	-0.00	0.03	-0.07	-0.05	-0.07
	(3.31)	(7.98)	(1.08)	(0.10)	(0.01)	(0.04)	(0.06)	(0.07)	(0.07)
Known Address	-0.26	7.09	0.45	0.04	-0.00	0.00	-0.08	-0.04	-0.03
	(3.29)	(7.35)	(1.03)	(0.09)	(0.01)	(0.04)	(0.06)	(0.07)	(0.07)
Observations	1303	1303	1303	1303	1303	975	422	341	341
Status Quo Average	62.50	52.43	8.32	0.46	0.03	0.42	0.40	0.64	0.63
Address Omission=KnownAddress p	0.57	0.81	0.85	0.22	0.99	0.51	0.86	0.90	0.53

Note: Sample includes only those who did not participate in the Information Experiment (as in the main text). Outcome in column: (1) whether the candidate clicked the link to start the application form in the WhatsApp message; (2) how many words applicants wrote in response to optional question at the end of the application form in which they could freely introduce themselves and say why they thought they were a good fit; (3) GPT-4 rating of the aforementioned response; (4) total years of experience declared in the application form; (5) whether any courses or experiences declared in the application form could be easily linked to a favela address or institution; (6) whether the participant responded to the endline *phone* survey (smaller sample size as not all batches were contacted); (7) whether the endline participant declared applying for another job besides the ones in this study; (8) whether jobseeker thought that the HR firm knew their address before sending the application invite; (9) whether jobseeker thought that the HR firm knew their phone number before sending the application invite.

Table B.2: Treatment Effects on Secondary Outcomes in the Address Omission Experiment, by Race

	(1) Clicked application link (%)	(2) Words in cover letter	(3) Cover letter quality (0/100)	(4) Years of experience declared	(5) Declared experience in favela (0/1)	(6) Participated in endline (0/1)	(7) Applied for another job (0/1)	(8) Thought HR knew address before applying (0/1)	(9) Thought HR knew phone before applying (0/1)
Address Omission	-5.94	0.93	-0.62	0.07	0.00	0.01	-0.09	-0.04	-0.06
	(3.80)	(9.61)	(1.24)	(0.11)	(0.01)	(0.05)	(0.07)	(0.08)	(0.08)
Known Address	-2.10	7.82	0.66	0.02	-0.00	0.00	-0.07	-0.01	-0.06
	(3.71)	(8.36)	(1.19)	(0.11)	(0.01)	(0.04)	(0.07)	(0.07)	(0.07)
Address Omission × white	15.86**	17.42	5.21**	0.36	-0.02	0.06	0.10	-0.01	-0.01
	(7.78)	(16.70)	(2.53)	(0.24)	(0.03)	(0.09)	(0.14)	(0.16)	(0.16)
Known Address × white	7.89	-4.60	-1.22	0.06	-0.02	-0.02	-0.04	-0.15	0.17
	(8.08)	(17.82)	(2.35)	(0.20)	(0.03)	(0.09)	(0.14)	(0.17)	(0.17)
Observations	1303	1303	1303	1303	1303	975	422	341	341
Status Quo Average	62.50	52.43	8.32	0.46	0.03	0.42	0.40	0.64	0.63

Notes: See the notes to the previous table.

Table B.3: Treatment Effects on Secondary Outcomes in the Information Experiment

	(1) Words in cover letter	(2) Cover letter quality (0/100)	(3) Exp. Disc. on the job (% chance)	(4) Exp. disc for self (%)	(5) Wage gap for self (%)	(6) # of planned future applications	(7) Exp. future callback rate (%)	(8) More excited about job search (0/1)	(9) Agrees job mkt is unfair (-2/2)	(10) Agrees job search is an opportunity (-2/2)	(11) Agrees one must ignore disc. (-2/2)	(12) Agrees one must apply to all jobs (-2/2)	(13) Plans to increase search effort (0/1)	(14) Feels is in bad position by endline (1/4)	(15) Improved search strategy by endline (0/1)
Favela Info	15.80	1.50	-6.47**	-2.77	-1.26	-8.27	-2.74	-0.22***	-0.01	-0.13	0.21*	-0.04	-0.02	-0.04	-0.10
	(10.34)	(1.57)	(2.54)	(3.09)	(2.95)	(6.46)	(2.66)	(0.04)	(0.12)	(0.11)	(0.11)	(0.13)	(0.03)	(0.11)	(0.06)
Full Info	15.41	0.33	-5.29**	-14.63**	-9.29	-3.37	0.72	-0.15***	-0.13	0.09	0.00	0.10	0.01	0.14	-0.04
	(10.36)	(1.54)	(2.66)	(6.62)	(8.59)	(6.92)	(2.78)	(0.04)	(0.13)	(0.11)	(0.11)	(0.13)	(0.03)	(0.11)	(0.06)
Observations	690	690	690	689	690	690	690	670	670	670	670	670	690	389	389
No Info Average	47.89	8.61	47.64	13.63	11.60	50.38	43.36	0.79	0.80	1.02	0.98	0.75	0.87	2.43	0.56
Favela=Full *p*	0.97	0.45	0.65	0.07	0.35	0.42	0.19	0.11	0.32	0.04	0.06	0.25	0.40	0.06	0.34

Note: Sample includes only those who did not participate in the Address Omission Experiment, as in the main text. Outcome in column: (1) how many words applicants wrote in response to optional question at the end of the application form in which they could freely introduce themselves and say why they thought they were a good fit; (2) GPT-4 rating of the aforementioned response; (3) predicted chance of suffering antifavela discrimination over the first year working in a formal job outside favela; (4) expected gap in employment probability in the next six months against someone similar living just outside the favela; (5) wage gap againt the same counterfactural person as in the previous column (a negative number implies antifavela wage discrimination); (6) number of applications the jobseekers wants to send in the next two months; (7) expected callback rate to the applications referred to in the previous column; (8) whether the jobseeker feels more excited about their job search at the end of the survey; (9) agreement with "the job market is extremely unfair" (Likert scale, -2=completely disagrees, 2=completely agrees); (10) agreement with "the job search is an opportunity to find the place I best fit into", same scale; (11) agreement with "to do well in the labor market, we can not think about employer discrimination all the time", same scale; (12) agreement with "to do well in the labor market, you have to apply to all possible vacancies", same scale; (13) whether one plans to increase their job search efforts over the next two months; (14) whether the endline survey respondent thinks that someone like them, from their neighborhood, has [NO=1/SOME=2/GOOD=3/GREAT=4] chance of finding a new formal job fast, (15) whether the endline survey respondent says they have worked on their résumé and took new measures to improve the odds they will find a job.

Table B.4: Treatment Effects on Secondary Outcomes in the Information Experiment, by Race

	(1) Words in cover letter	(2) Cover letter quality (0/100)	(3) Exp. Disc. on the job (% chance)	(4) Exp. disc for self (%)	(5) Wage gap for self (%)	(6) # of planned future applications	(7) Exp. future callback rate (%)	(8) More excited about job search (0/1)	(9) Agrees job mkt is unfair (-2/2)	(10) Agrees job search is an opportunity (-2/2)	(11) Agrees one must ignore disc. (-2/2)	(12) Agrees one must apply to all jobs (-2/2)	(13) Plans to increase search effort (0/1)	(14) Feels is in bad position by endline (1/4)	(15) Improved search strategy by endline (0/1)
Favela Info	14.37	1.30	-3.82	-2.61	0.87	-8.11	-4.33	-0.23***	0.09	-0.09	0.25**	-0.06	-0.02	-0.08	-0.06
	(11.81)	(1.73)	(2.92)	(2.92)	(3.31)	(7.71)	(2.98)	(0.05)	(0.14)	(0.12)	(0.12)	(0.14)	(0.04)	(0.13)	(0.07)
Full Info	18.26	0.92	-4.44	-19.55**	-9.80	-10.86	-1.89	-0.14***	-0.13	-0.04	-0.04	0.06	-0.01	0.12	0.02
	(11.72)	(1.73)	(3.04)	(8.04)	(10.87)	(7.33)	(3.16)	(0.05)	(0.14)	(0.12)	(0.13)	(0.14)	(0.04)	(0.12)	(0.07)
Level Info × white	6.85	0.96	-12.82**	-0.72	-10.45	-0.53	7.72	0.03	-0.45	-0.19	-0.19	0.07	-0.00	0.20	-0.19
	(24.61)	(4.14)	(5.79)	(10.56)	(7.28)	(12.33)	(6.68)	(0.11)	(0.29)	(0.29)	(0.28)	(0.33)	(0.08)	(0.27)	(0.15)
Full Info × white	-13.27	-2.77	-3.97	22.55*	1.62	35.39*	12.07*	-0.02	-0.04	0.55**	0.20	0.19	0.09	0.11	-0.32**
	(25.52)	(3.89)	(6.30)	(12.95)	(12.78)	(19.01)	(6.65)	(0.11)	(0.30)	(0.25)	(0.26)	(0.32)	(0.07)	(0.26)	(0.16)
Observations	690	690	690	689	690	690	690	670	670	670	670	670	690	389	389
No Info Average	47.89	8.61	47.64	13.63	11.60	50.38	43.36	0.79	0.80	1.02	0.98	0.75	0.87	2.43	0.56

Notes: See the notes to the previous table.

Table B.5: Secondary Outcomes in the Interview Experiment

	(1) Nervousness cues (0/1)	(2) Gave away address (0/1)	(3) Used slang (0/1)	(4) Aggregate question-wise performance (SD)	(5) Interviewer professionalism (perceived, SD)	(6) Interviewer preparedness (perceived, SD)	(7) HR firm values diversity (SD)
Name-Only	0.00	0.03	-0.01	0.08	-0.07	0.07	0.07
	(0.03)	(0.02)	(0.01)	(0.08)	(0.11)	(0.09)	(0.08)
Observations	422	422	422	422	422	422	422

Notes: Outcome in column: (1) whether the interviewer noted that the candidate laughed out of nervousness, stuttered, or had a shaking voice; (2) whether the candidate explicitly gave away their neighborhood during the interview; (3) whether the candidate used slang during the itnerview; (4) ICW index of performance in the six main interview questions; (5) from the candidate's feedback form, normalized rating of the interviewer's professional behavior; (6) from the candidate's feedback form, normalized rating of how prepared the interviewer seemed to be; (7) from the candidate's feedback form, normalized rating of how much it seemed like the HR firm valued diversity. Outcomes in columns (5) to (8) had little variation, about 85% of candidates picked ten out of ten in those questions.

Table B.6: Secondary Outcomes in the Interview Experiment

	(1) Nervousness cues (0/1)	(2) Gave away address (0/1)	(3) Used slang (0/1)	(4) Aggregate question-wise performance (SD)	(5) Interviewer professionalism (perceived, SD)	(6) Interviewer preparedness (perceived, SD)	(7) HR firm values diversity (SD)
Name-Only	0.01	0.02	-0.01	0.02	-0.11	0.06	0.08
	(0.04)	(0.02)	(0.01)	(0.09)	(0.12)	(0.10)	(0.10)
Name-Only × white	-0.03	0.06	0.03	0.25	0.17	0.03	-0.03
	(0.06)	(0.04)	(0.02)	(0.18)	(0.27)	(0.22)	(0.20)
Observations	422	422	422	422	422	422	422

Notes: See the notes to previous table.

B.2 Main Tables and Figures Including Individuals Who Participated in Both the Address Omission and Information Experiments

Tables and figures related exclusively to the interview experiment or already including the full sample are omitted, since they would not change.

Table B.7: Table 1 Including Sample Overlapping with Address Omission Experiment

	(1) Responded to endline (0/1)	(2) Exp. discrimination (categorical, 1-4)	(3) # of sent apps (categorical, 1-4)
Favela Info	-0.00	0.02	-0.03
	(0.04)	(0.08)	(0.12)
Full Info	0.01	-0.15*	0.06
	(0.04)	(0.09)	(0.12)
Observations	864	506	508
Controls			
No Info Mean	0.6	2.3	2.5
Favela=Full *p*	0.78	0.04	0.45

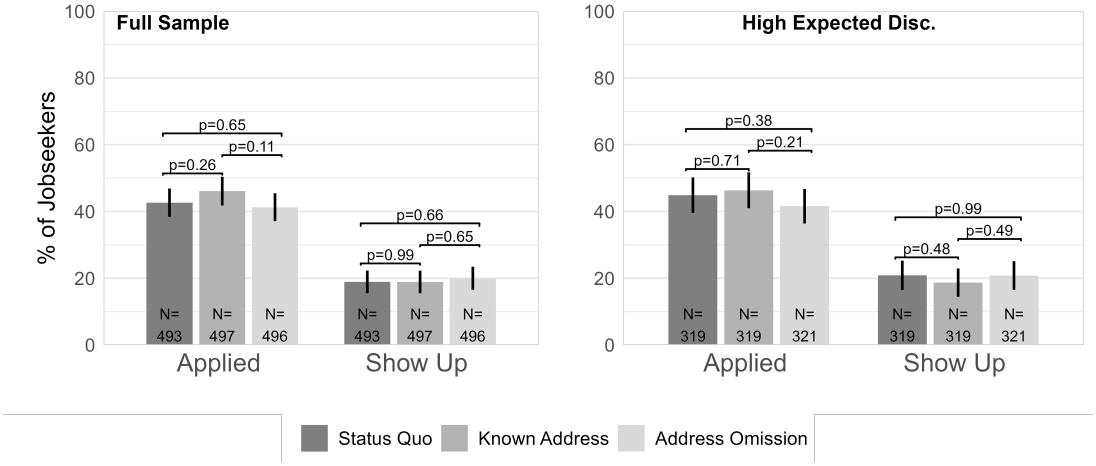
Note: See notes for Table 1 for details.

Table B.8: Table 3 Including Sample Overlapping with Address Omission Experiment

	Experiment-specific outcomes			
	(1)	(2)	(3)	(4)
Agg. interv. performance (SD) *Name-Only* $\times h_i$	0.24*	0.24**	0.07	−0.12
	(0.14)	(0.12)	(0.13)	(0.12)
Show-up (%) *Address Omission* $\times h_i$	5.56	−3.16	−2.64	0.71
	(6.28)	(5.12)	(5.21)	(5.04)
Show-up (%) *Known Address* $\times h_i$	−9.85*	−6.50	9.22*	−5.89
	(5.86)	(5.15)	(5.57)	(5.02)
Show-up (%) *Full Info* $\times h_i$	−13.71*	−2.70	−8.29	9.32
	(7.93)	(6.93)	(7.40)	(6.85)
Show-up (%) *Favela Info* $\times h_i$	−1.70	1.79	−5.86	9.66
	(8.27)	(6.76)	(7.16)	(6.63)
Heterogeneity variable h	White	High $\mathbb{E}$[disc]	Male	High Skill
Any heterogeneity by h, p-value	0.021	0.286	0.252	0.299
Heterogeneity by h in $\mathbb{E}$[address visib.] treatments, p-value	0.017	0.117	0.156	0.358
Heterogeneity by h in information treatment, p-value	0.161	0.809	0.51	0.264
Clusters	2167	2167	2167	2167
Observations	2,772	2,772	2,772	2,772

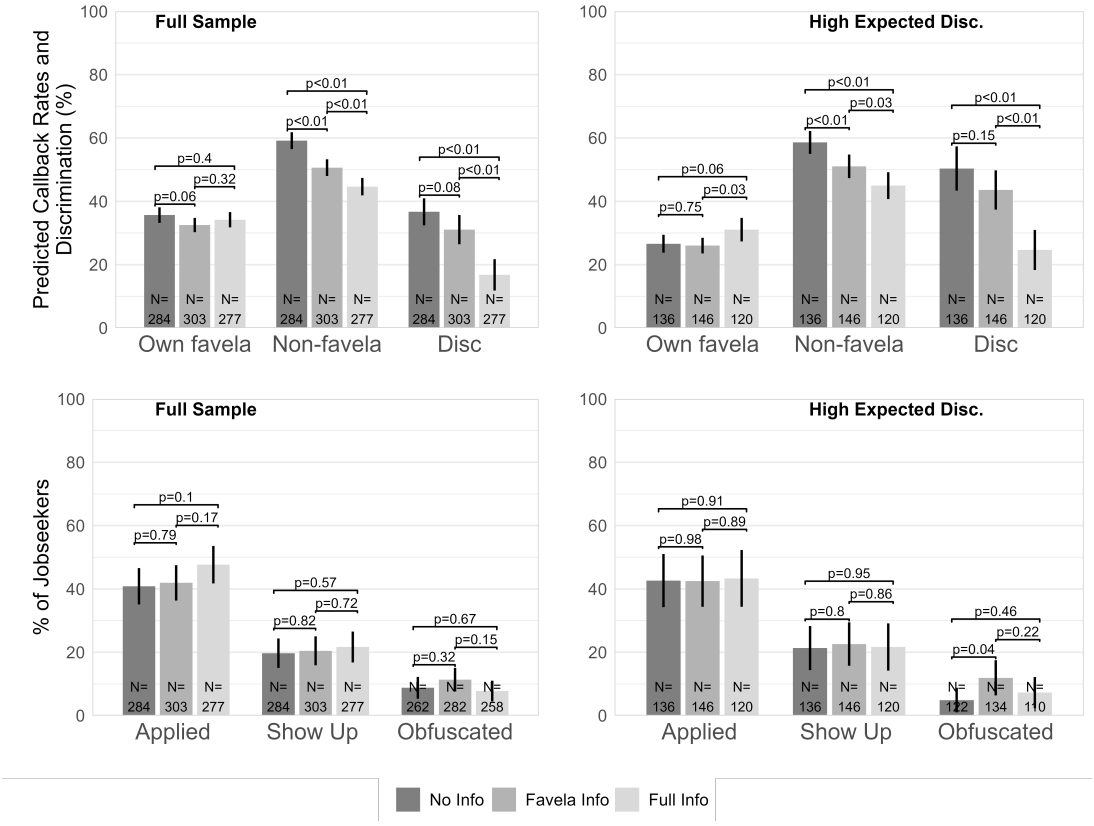
Note: See notes to Table 3.

Figure B.1: Figure 4 Including Sample Overlapping with Information Experiment



Note: See notes to Figure 4 for details.

Figure B.2: Figure 5 Including Sample Overlapping with Address Omission Experiment



Note: See notes to Figure 5 for details.

C Audit Study

Picking common names. We picked random combinations of first, middle and last names among the 50 most common possibilities (by gender) among all formal workers in the state of Rio de Janeiro (from the RAIS dataset). In Brazil, some names may be distinctive of lower socioeconomic status, but names that are distinctive in terms of race are very rare. At any rate, the selected names were so common that they were not distinctive in any way.

Résumé addresses. For addresses in each neighborhood, we picked streets that were i) entirely contained in the neighborhood, ii) in the postal office list for that neighborhood, 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.

Selecting vacancies. We found vacancies through 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.

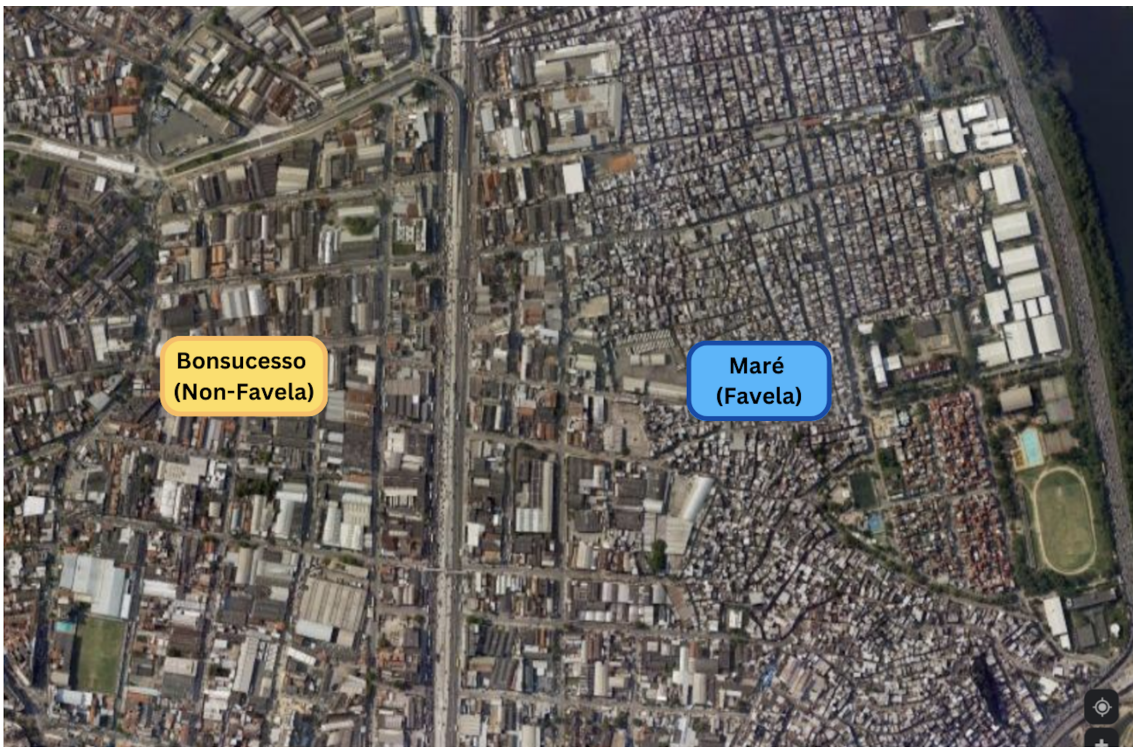
Randomization. The exact randomization procedure was that, for each job posting, we first randomly ordered the four profiles. Among the first two randomly ordered profiles, we randomly picked one for being from Maré. We did the same for the latter two, which were backups. A research assistant applied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

Table C.1: Audit Study Results

	(1) Callback (%)	(2) Callback (%)	(3) Callback (%)	(4) Callback (%)
Maré résumé	-0.34	-0.39	0.66	-1.04
	(1.28)	(1.29)	(2.08)	(1.18)
Maré × Downtown			-2.45	
			(3.99)	
Maré × South			2.64	
			(4.19)	
Maré × Further south			-3.34	
			(3.37)	
Observations	1400	1400	1174	1400
Non-favela Mean	16.96	16.96	16.75	16.96
Controls	No	Yes	Yes	No
Job FEs	No	No	No	Yes

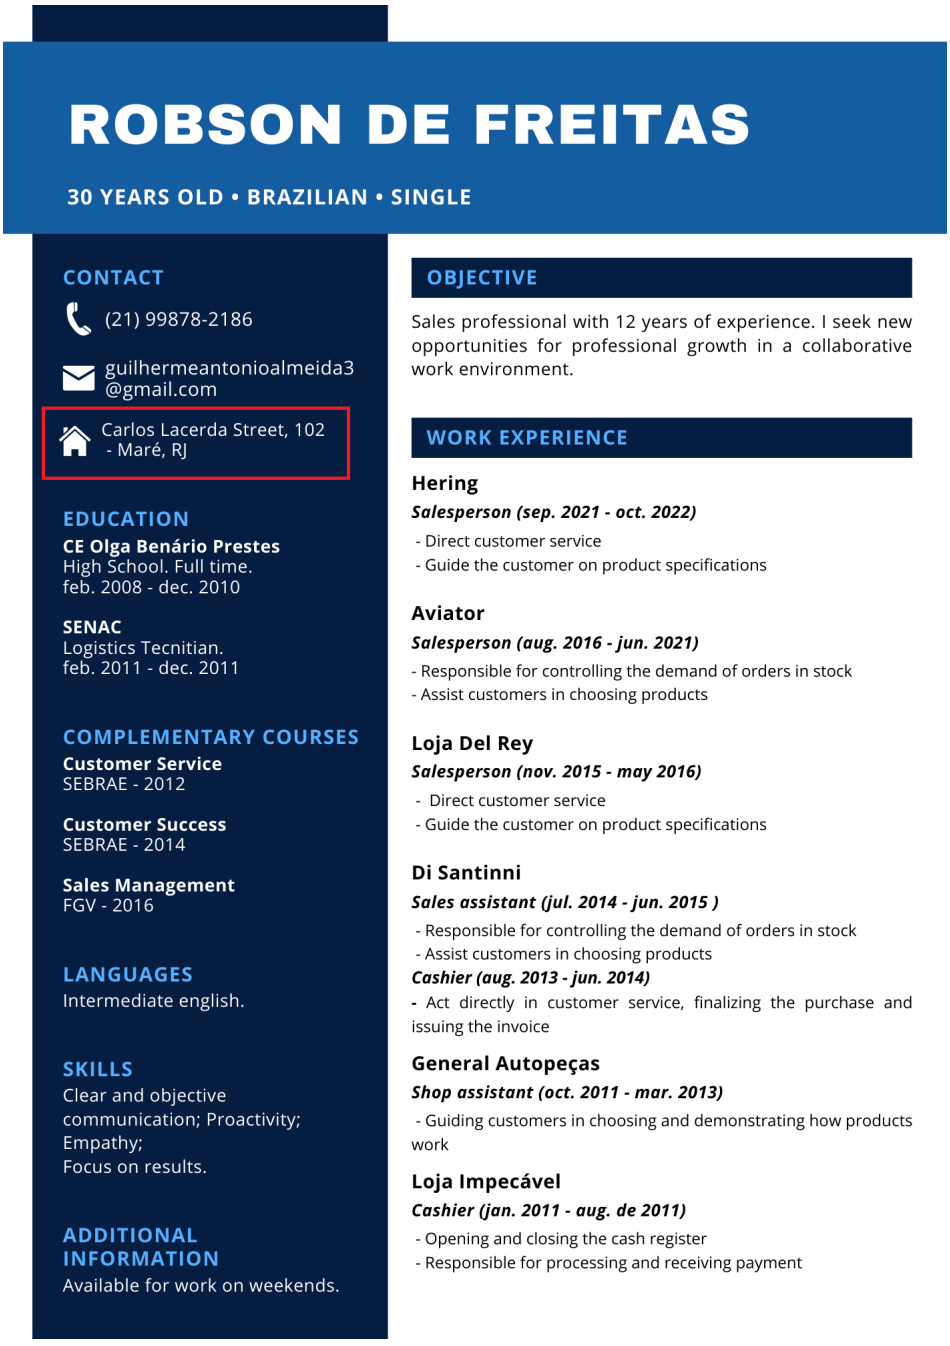
Note: Outcome variable evaluates to 100 if the application received a positive response and zero otherwise. Maré résumé is a dummy for the fictitious applicant being from Maré. Controls include the neighborhood in which the job is located, the website that advertised the job, dummies for the number of vacancies in the ad, and the hiring regime. In column 3, the sample is restricted to jobs that coders could trace back to a single location out of four possibilities: (i) Zona Norte (the omitted category), which includes Maré and Bonsucesso, (ii) Downtown, where we had our interview office, (iii) Zona Sul (South), which is considerably richer than Zona Norte, and (iii) Barra da Tijuca ou Jacarépagua (Further south), which is also richer but further away. The callback level in this table is $\approx$ 3p.p. lower than in the main text because the regressions models only consider callbacks we could link to unique postings. Standard errors clustered at the posting level shown between parenthesis.

Figure C.1: Sattelite Image: Bonsucesso (Nonfavela) and Maré (Favela)



Note: The large avenue (vertical) in the picture is the separates the neighborhoods.

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

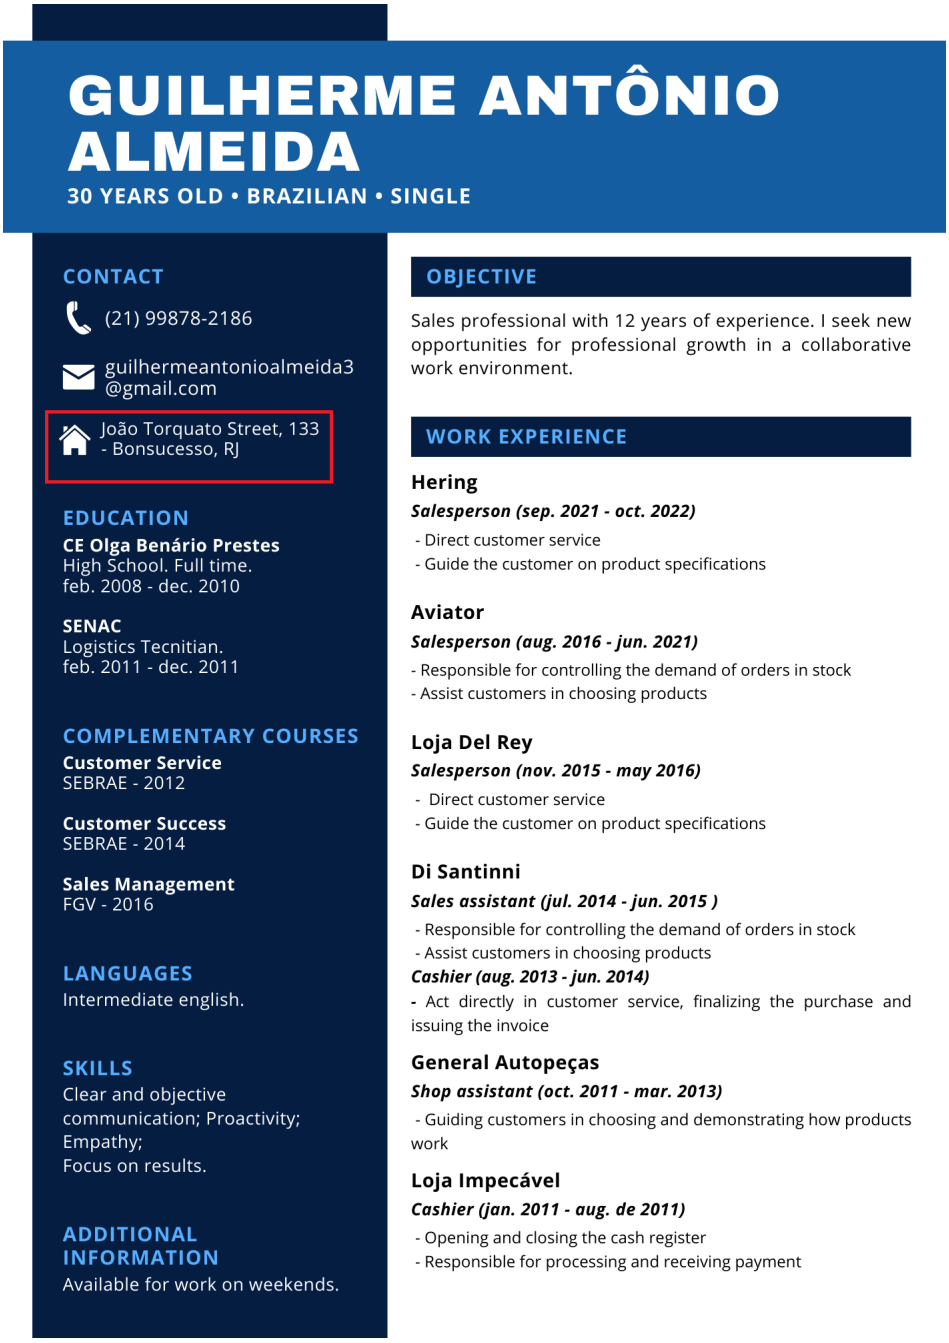
Cashier (aug. 2013 - jun. 2014)
- Act directly in customer service, finalizing the purchase and issuing the invoice

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. The red box around the address in this picture was added for emphasis, it was not present in the original résumé.

Figure C.3: Example Résumé – Bonsucesso Address



Note: This image shows nonfavela counterpart of the résumé above.

D Supporting Materials

This section includes the key scripts and forms used in the study, as well as photos illustrating the procedures.

D.1 Door-to-door Survey

D.1.1 Survey Excerpts

Introducing the HR firm. After determining eligibility, going over informed consent, and asking questions about demographics and job market experience, the surveyor read the following italicized text:

> *The organizing team behind this study has a partnership with S.A.M. RH, a recruitment company here in Rio, that helps large companies find employees.*
>
> *If you authorize, we can send your basic profile, including the responses you've given so far, to S.A.M. RH, and they may contact you to apply for some jobs, if you meet the prerequisites.*
>
> *Do you authorize us to share your information with S.A.M. RH?*

Expected callback elicitation. The script below was used for surveys in Maré. Surveys completed in other favelas also included elicitations of the callback rates for that other favela (and Maria da Graça, in the case of Jacarézinho). The description of the audit study did not mention specific neighborhoods in surveys in favelas other than Maré.

> *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 R$100.*
>
> *Let me tell you the whole story. At the beginning of our project, the researchers organizing this study heard from the population of several favelas here in Rio about how much harder it was to apply for a formal job for someome living in a community. To really understand the size of the challenge, researchers sent 1,400 applications with fake résumés, but as if they were real people, for 700 vacancies in sales in the city of Rio.*
>
> *The résumés were from men and women, from people with experience and suitable for each vacancy. The only difference between the résumés 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 the sent résumés.

[Figure C.2 appeared here.]

The researchers calculated WHAT PERCENTAGE of résumés sent with BONSUCESSO's address were selected (for example, for a training period) or invited for an interview. They also calculated this percentage for MARÉs résumés.

For the money prize, I'm going to ask you to guess what they found, okay?

- *WHAT PERCENTAGE of résumés with BONSUCESSO's address do you guess were selected or invited for an interview?*
- *AND WHAT PERCENTAGE of MARÉ's?*



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

D.2 Application Form

Figure D.2: 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

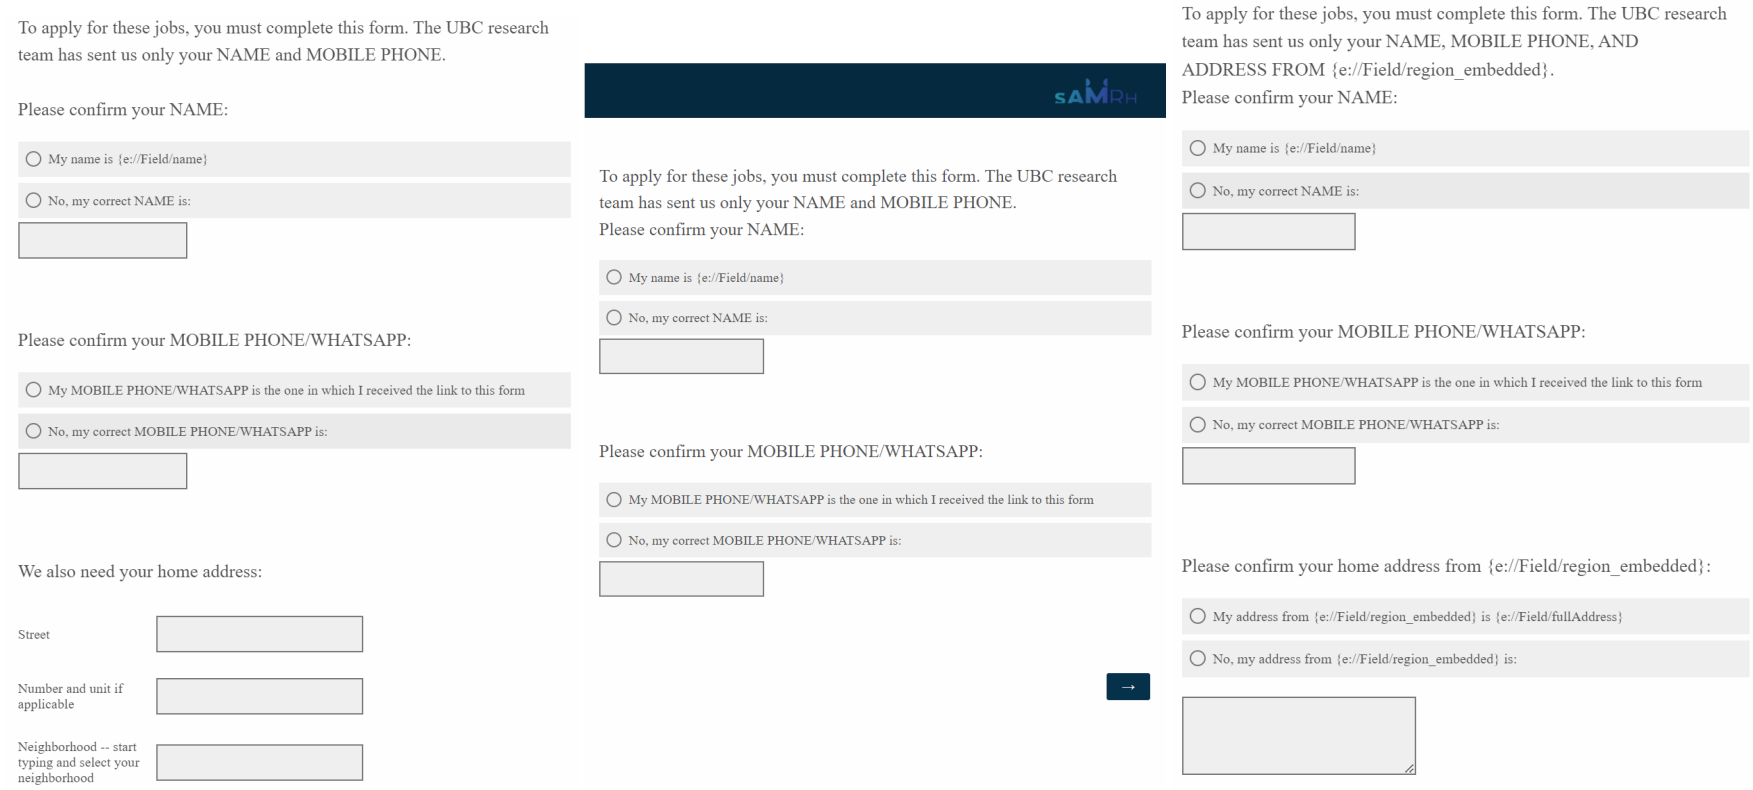
Duties:
- Responsible for receiving payments, operating sales systems, issuing invoices, making contact with resellers

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 (translated from Portuguese).

Figure D.3: Second Screen of the Application Form of Each Experimental Condition in the Address Omission Experiment



(a) *Status Quo* (b) *Address Omission* (c) *Known Address*

D.3 Reception Script

The receptionist was in charge of scheduling logistics and directing the candidates inside the office. She had access to candidates' names, phone numbers, and date of birth, but not address. Hence, when the receptionists asks to confirm address (see below), she simply takes whatever candidates say at face value.

When a candidate arrived at her desk, she would follow this script (a Qualtrics form):

"Hello, how are you? I am [NAME], the receptionist here at SAM HR. Can you please confirm some information?"

Q1) Name:

- [Name provided in the application]
- Corrected name: _____

Q2) Date of birth:

- [DOB provided in the application]
- Corrected DOB: _____

Q3) Address:

Street: _____

Number: _____

Neighborhood: [Pick from drop-down list]

Ask the candidate to wait until interviewer is ready. When ready, or after a moment:

"Ok! Your interviewer today is [INTERVIEWER NAME]. Here at SAM HR we try to be very objective in our selection procedures, to pick the best candidates, so, because of that, she **will** only know your **[name/name and address], and nothing more about you**, ok?"

D.4 Interview Script

The italicized text below was *not* read out loud.

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 BOX

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 BOX
Interviewer evaluates how well the candidate did on this question, from 0 to 10

Q4. Now I will ask you to do an exercise. Think of a product you like and know well. It could be 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 BOX
Interviewer evaluates how well the candidate did on this question, from 0 to 10, and notes: (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 BOX
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 BOX
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 need to talk about professional experiences, necessarily. It could be something academic, school-related, 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 BOX

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?*

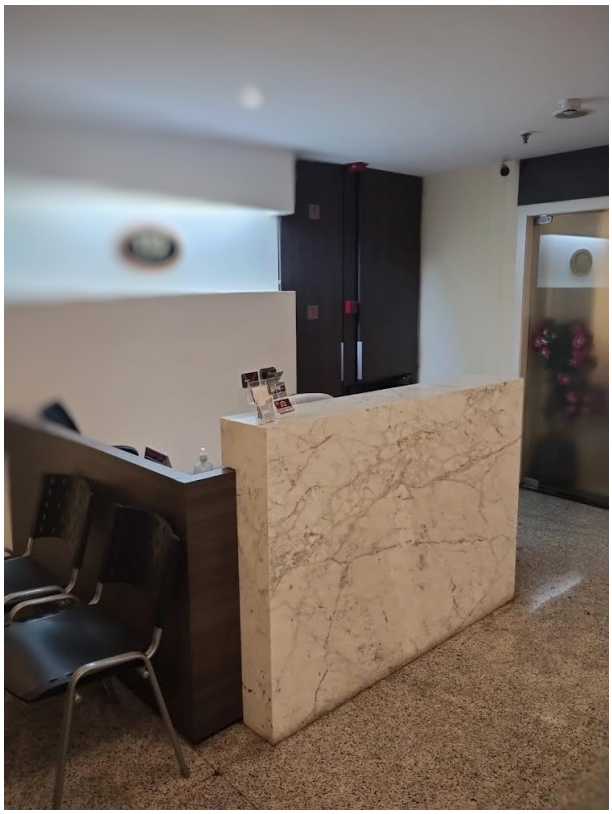
2. *How nervous did the candidate seem?*
3. *How focused did the candidate seem?*
4. *How professional was the candidate throughout the interview?*

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

- *Had a shaky voice*
- *Stuttered*
- *Laugh nervously*
- *Dressed in informal clothes*
- *Used slangs*
- *Made MANY Portuguese language mistakes*
- *Used swear words*
- *Mentioned personal things, irrelevant to the position*
- *Mentioned that they were religious or went to church or worship*
- *Mentioned that they lived in a favela*
- *Talked about where they came from (on that day)*
- *Talked about where they lived*
- *Talked about where they were born*
- *Asked you personal questions*
- *Asked you irrelevant questions for the position*
- *Showed you know they knew something(s) about the company or the position*
- *Used very formal language*
- *Looked you in the eyes when answering*
- *Avoided looking into your eyes*
- *Was very shy*
- *None of the above*

Figure D.4: Spaces Used in the Interview Experiment

(a) Reception



(b) Interview Room

