

Expected Discrimination and Job Search*

Deivis Angeli[†]

Ieda Matavelli[‡]

Fernando Secco[§]

August 9, 2026

Abstract

We study how expected discrimination affects job applications and interview performance in three field experiments with 2,167 jobseekers living in Brazilian favelas (urban slums). We focus on antifavela discrimination, which 87% of jobseekers overestimate. Randomizing expected address visibility—or providing information about discrimination in callbacks—does not affect average application rates or interview attendance. However, expecting interviewers to know one’s favela address reduces interview performance by 0.13SD, even though interviewers are in fact blind to addresses. Expected discrimination can thus affect labor market matching, especially in hiring processes that involve face-to-face interviews.

*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, Karla Hoff, Supreet Kaur, Ro’ee Levy, Federico Masera, Nathan Nunn, Devin Pope, Gautam Rao, Chris Roth, Rogério Santarrosa, Heather Sarsons, Colin Sullivan, Jonathan Zinman, seminar participants at the Vancouver School of Economics, and three anonymous referees 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-0011041). All data and code to reproduce the results in this article are available from the AEA Data and Code Repository at <https://doi.org/10.3886/E251550V1>. The views expressed in this article are those of the authors and do not necessarily reflect those of Analysis Group or its clients.

[†]Global Talent Lab, deivisangeli@gmail.com

[‡]Monash University; 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—their beliefs about receiving unfair or prejudicial treatment based on characteristics they possess—remain understudied, despite their theoretical importance for labor market outcomes. If jobseekers overestimate discrimination, they may avoid applying to jobs or perform poorly in interviews due to stress, even when employers do not discriminate. Such reactions could amplify the effects of employer discrimination and create self-reinforcing loops (Coate and Loury, 1993). Understanding when and how expected discrimination affects job search therefore also matters for policymakers and firms: Correcting misperceptions could prevent self-fulfilling prophecies, and firms could attract more applicants by credibly signaling procedural fairness.

This paper studies how expected discrimination affects job application decisions and interview performance with three field experiments. We focus on address-based discrimination as perceived by favela residents in Rio de Janeiro, Brazil, where 1.5 million people (22% of the population) live in favelas. Unlike characteristics such as race, address is not inherently visible during hiring processes since it must be actively disclosed, which allows us to randomize expected address discrimination by manipulating expected address visibility. Our experiments follow a sequential pipeline designed to mirror natural job search processes while maintaining experimental control. We first recruit favela jobseekers through door-to-door surveys in three favelas in Rio. A few days later, an HR firm (which we operate) invites participants to apply for real full-time sales jobs with a partner company. We run two experiments at this application stage: One varying the extent to which jobseekers expect the firm to know their addresses and one varying whether we provide accurate information about market-level discrimination. We then invite those who completed applications to attend a job interview, where we manipulate candidates’ expectations about whether interviewers would know their addresses.

We begin by documenting that most favela jobseekers overestimate antifavela discrimination. In a door-to-door survey with 2,167 favela residents, we elicited incentivized predictions of the callback rates we would find in an audit study we were conducting simultaneously. For the audit study, we created fictitious résumés that varied by whether the listed address was in a favela or in an adjacent nonfavela neighborhood, and then submitted those résumés to 700 sales jobs. In our survey of jobseekers, 87% predicted that résumés with favela addresses would receive fewer callbacks than those with nonfavela addresses. The median jobseeker predicted that a favela address would cause the callback rate to drop by 50%, while our audit study finds sta-

¹See 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.

tistically indistinguishable callback rates of 19.3% for favela résumés and 19.6% for nonfavela résumés. The small difference in callback rates does not imply an absence of discrimination, since discrimination may occur at later stages in the hiring process. The audit study provides an objective benchmark, revealing that jobseekers are too pessimistic about antifavela discrimination at the callback stage.

Our first experiment tests whether expected address visibility affects job application decisions and interview attendance (Address Omission Experiment; N=1,303). We randomly assigned favela jobseekers to one of three conditions when inviting them to apply for real sales jobs. Application forms either requested an address (*Status Quo*), omitted the address field entirely (*Address Omission*), or listed a pre-filled favela address (*Known Address*). *Status Quo* emulates the typical practice of asking for an address, giving us a natural point of comparison and allowing us to measure whether jobseekers behave strategically by providing nonfavela addresses to avoid discrimination. Comparing *Known Address* and *Address Omission* allows us to measure the effect of expected address visibility to its full extent, since *Known Address* guarantees that jobseekers expect their favela addresses to be known while *Address Omission* guarantees that they expect addresses to remain hidden.

Despite this experimental manipulation, we find no average effect of expected address visibility on application rates or interview attendance. Application rates are between 42-45% across the three conditions, and interview attendance rates are 17-20%, with no statistically significant differences between any pair of conditions. These null effects persist even among jobseekers who predicted discrimination rates of 50% or above in the audit study belief elicitation. A failure to manipulate expected address visibility cannot fully explain these null average effects: In *Status Quo*, 44% of applicants obfuscated their addresses by declaring a different neighborhood, demonstrating that jobseekers indeed perceived address-based discrimination and responded strategically when possible. Although average effects are null, we find suggestive evidence that white jobseekers become more likely to attend the interview as address visibility decreases, while nonwhite jobseekers seem unaffected. This pattern is consistent with a “passing” mechanism: Since only 33% of favela residents are white (versus 57% of nonfavela residents), white jobseekers can “pass” as nonfavela residents when addresses remain hidden, avoiding discrimination based on characteristics associated with favela residence. Nonwhite jobseekers, however, may still expect racial discrimination or to be stereotyped as favela residents, regardless of address visibility.

Next, our second experiment tests whether expected discrimination hurts interview performance (Interview Experiment; N=422). We invited all jobseekers who completed applications to attend real job interviews at a downtown office. Upon arrival, a receptionist collected can-

didates' names and addresses, then randomly assigned them to one of two conditions: (i) *Name-Only*, where the receptionist stated that “the interviewer will only know your name” or (ii) *Name-and-Address*, where the receptionist stated that “the interviewer will only know your name and address”. Importantly, interviewers were completely blind to candidates' addresses in both conditions, knowing only candidates' names. The treatment therefore operated purely through candidates' expectations about what interviewers knew, allowing us to identify the causal effect of expected discrimination independent of actual discriminatory evaluation. Immediately after each interview, interviewers rated candidates on overall performance, nervousness (reverse-coded as calmness), and professionalism, all on 0-10 scales. Candidates also completed private self-assessments on the same three dimensions. Based on these measures, we construct inverse-covariance-weighted indices of interviewer-assessed and self-assessed performance (Anderson, 2008), then average both to create an aggregate performance index.

Expecting interviewers to know their addresses hurts interview performance. The *Name-and-Address* condition decreases aggregate performance by 0.13SD, compared to *Name-Only*. This aggregate effect is driven primarily by effects on self-assessments (an effect of 0.17SD, with $p < 0.01$) and less so by effects on interviewer assessments (0.09SD, $p = 0.28$), though we cannot reject that the two effects are equal ($p = 0.33$). All six individual performance components move in the same direction, with effects varying from 0.06SD to 0.15SD. The interview performance effects concentrate among jobseekers who expect high discrimination (0.22SD, $p < 0.05$) and white jobseekers (0.3SD, $p < 0.01$). The heterogeneity by expected discrimination is consistent with a self-fulfilling prophecy: Candidates who expect substantial discrimination perform worse when they believe interviewers know their addresses, even though interviewers remain blind. This suggests that even if employer discrimination is absent, gaps in outcomes could persist through jobseekers' pessimistic beliefs. Consistent with the “passing” pattern observed in the Address Omission Experiment, we find suggestive evidence that expecting interviewers to know their addresses hurts performance particularly among white jobseekers.

The Address Omission and Interview Experiments manipulate expected address visibility (and, consequently, expected discrimination) within a specific job application procedure. In our third experiment, we test whether providing information about market-level discrimination shifts beliefs about expected discrimination and encourages job applications (Information Experiment; $N = 690$). We randomly assigned survey respondents to one of three conditions: (i) *No Info*, where we revealed no information about our audit study results; (ii) *Favela Info*, where we revealed only the favela callback rate (19.3%); and (iii) *Full Info*, where we revealed both favela and nonfavela callback rates (19.3% and 19.6%).

The information treatments successfully shifted beliefs, but did not affect application behav-

iors and interview attendance. Learning our audit results reduced average posterior expected discrimination from 36% in *No Info* to 29% in *Favela Info*, and 17% in *Full Info*, with differences statistically significant. However, despite these substantial belief shifts, we find no statistically significant differences in application rates or interview attendance across conditions. Application rates range from 40-45% and interview attendance rates from 19-21% across the three conditions, with no statistically significant differences across arms. Even among jobseekers who initially expected high discrimination, the information treatments do not change application rates.

Taken together, our main finding is that expected discrimination significantly impairs interview performance but does not affect average application or interview attendance rates, despite jobseekers’ overestimation of discrimination. In addition, providing market-level information on actual discrimination does not increase application rates. While we cannot pin down the exact mechanisms through which expected discrimination affects interview performance, one possibility is that jobseekers choke under pressure. The null effect of information on application decisions could be explained by jobseekers finding market-level statistics insufficiently persuasive to change behavior. We discuss these and alternative mechanisms in Section 5.1.

This paper contributes to several strands of the literature. First, we contribute to an extensive body of empirical work on labor market discrimination. While many experiments measure whether employers discriminate (Riach and Rich, 2002; Rich, 2014; Bertrand and Duflo, 2017; Neumark, 2018), jobseekers’ reactions to expected discrimination have only recently received more attention. Lab and online experiments show that people strategically respond to expected discrimination by crafting “whitened” résumés (Kang et al., 2016) and hiding their gender (Alston, 2019; Charness et al., 2020). While these experiments demonstrate strategic concealment, they tell us little about how expected discrimination affects downstream behaviors. Exceptions that go beyond strategic concealment find that gender-blind résumés encourage women to apply (Boring et al., 2025) and that women prefer algorithmic rather than male evaluators (Pethig and Kroenung, 2023). We add to this work by providing field evidence on how expected discrimination affects actual application and interview behavior with real jobs and high-stakes. Additionally, by focusing on address-based discrimination—a non-visible characteristic—we can study how expected discrimination impacts performance in face-to-face interactions, which lab studies have generally not examined because they often center on always-visible characteristics.

In field settings, our paper relates to experiments manipulating recruitment procedures. These studies randomize job ad language (Del Carpio and Fujiwara, 2025; Burn et al., 2022), photographs of current workers (Delfino, 2024), diversity or affirmative action statements (Flory

et al., 2022; Leibbrandt and List, 2025; Ibañez and Riener, 2018), and the use of AI evaluators (Avery et al., 2024) to study labor supply responses. While related to expected discrimination, these papers do not provide evidence that expected discrimination itself is the mechanism driving their results. Avery et al. (2024) is an exception, but measures expected discrimination only after treatment and among jobseekers who completed their job applications, making interpretation difficult. Moreover, the manipulations of job application procedures used in these papers can affect applications for reasons other than expected discrimination. For instance, gender-neutral job ads may encourage women not only by reducing expected discrimination but also by signaling better work-life balance or different job requirements (Del Carpio and Fujiwara, 2025). Similarly, diversity statements may signal organizational values beyond discrimination. Our experiments varying expected address visibility at the application stage remain subject to similar challenges. Nevertheless, not only are we able to measure whether baseline expected discrimination predicts treatment effects (in accordance with the hypothesized mechanism), we also complement this evidence with our Information Experiment, which holds recruitment procedures constant while directly shifting expected discrimination. Furthermore, we also integrate several aspects which have only been studied separately—strategic responses, applications, interviews, and belief elicitation—within one naturalistic setting. Related, but outside the realm of the job search, Ruebeck (2024) studies how expected discrimination affects on-the-job outcomes like retention and performance.²

Second, we contribute to a literature about the role of workers’ beliefs in determining labor market outcomes. Even without differences in initial endowments, pessimistic beliefs about returns to investment can make workers acquire less human capital in response to expected discrimination, such that discrimination becomes self-fulfilling (Coate and Loury, 1993; Lundberg and Startz, 1983).³ Glover et al. (2017) shows that minority employees working under biased

²Several observational studies find evidence consistent with expected discrimination affecting job search. For instance, Pager and Pedulla (2015) shows that Black jobseekers, especially those who have experienced discrimination, cast wider nets; Glover (2019) finds workers with Arabic-sounding names reduce search intensity following terrorist attacks; Kuhn and Shen (2023) shows that gender representation increases when employers remove explicit gender requests from job ads; and Agüero et al. (2023) suggests that jobseekers do not reveal a strong productivity signal when that signal would also reveal disadvantaged group status. Beyond job search, people appear to respond to expected discrimination by generating clearer productivity signals (Lepage et al., 2022; Dickerson et al., 2024; Lang and Manove, 2011) or changing major choices (Lepage et al., 2025). This literature typically proxies expected discrimination using past discrimination experiences, salience shocks, or non-incentivized self-reports, which makes it difficult to cleanly identify the role of expected discrimination from other correlated channels, even when quasi-experimental variation is available (e.g., Kuhn and Shen 2023).

³This work features “rational expectations” and statistical discrimination, exclusively. Goldsmith et al. (2004) later illustrated how another psychological mechanism (cognitive dissonance) may also cause expected discrimination to affect job search. Stereotype threat, a concept different from but related to expected discrimination, can also lead to similar self-fulfilling prophecies (see a discussion in Section 2.2). Outside labor market settings, Hoff and Pandey (2006, 2014) show how these self-fulfilling prophecies can stem from a stigma that represents a group

managers perform worse, creating a self-fulfilling prophecy that confirms managers’ biases, and [Bohren et al. \(2025\)](#) shows how expected *direct* discrimination can compound into *systemic* discrimination. We contribute by showing, in a field setting, that expected discrimination generates a self-fulfilling prophecy in the matching process through jobseekers’ expectations alone: Even when interviewers remain blind to addresses, jobseekers who expect their addresses to be known perform worse. Further, we also document a pervasive pessimism in beliefs about discrimination while replicating findings of excessive optimism about job-finding rates ([Mueller et al., 2021](#); [Bandiera et al., 2025](#); [Abebe et al., 2025](#)).

Third, we join a short list of papers in economics investigating the later stages of the job selection process.⁴ [Amer et al. \(2024\)](#) and [Shukla \(2024\)](#) find that interviews have a crucial role in creating gender and caste gaps. [Goldin and Rouse \(2000\)](#) shows that blind auditions increase female hiring in orchestras, which could reflect both reduced evaluator discrimination and improved female performance under blind evaluation. We add to this literature not only by directly measuring interview performance but also by showing that expected discrimination hurts performance. To our knowledge, [Godlonton \(2020\)](#) is the only other paper measuring face-to-face performance.

2 Expected Discrimination in Context

2.1 Favelas in Rio de Janeiro

Brazilian favelas are urban neighborhoods of dense informal settlements. These neighborhoods have emerged and persisted because they provide access to urban labor markets at substantially lower housing costs while simultaneously limiting further social mobility (e.g., due to lack of property rights and good schools, see [Ferreira et al. 2025](#)). In Rio de Janeiro, favelas are home to 1.5 million people, about one-fifth of the city’s population. According to the 2010 census (Table A.1; [Instituto Brasileiro de Geografia e Estatística, 2011](#)), 67% of favela households had a per capita income that was strictly positive but at most one minimum wage (≈ 10 USD/day). Outside the favela, this rate is 31%, and per capita income is 3.6 times larger. Favela residents are also less likely to be literate (84% are literate inside favelas, 92% outside them), and to self-identify as white (33% in favelas vs. 57% outside).

as inferior. [Hoff and Stiglitz \(2010\)](#) shows how modern history empires have created the beliefs to support these.

⁴Early audit studies, primarily by sociologists, hired actors from different groups to apply *and interview* for jobs while acting the same. Aggregate evidence from these studies suggests that most of the discrimination happens at the job interview stage rather than before it ([Quillian et al., 2020](#)). They also reveal nuance: [Pager et al. \(2009\)](#) shows how, when face-to-face, employers channel nonwhite candidates into worse jobs than those originally posted.

The jobseekers in our study lived in Maré, Mangueiras, 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.1 for census statistics comparing these neighborhoods). We conducted most of our fieldwork in Maré, the most populous favela in Rio.

Favelas offer limited formal work opportunities. 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. Formal 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, usually 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 the favelas in our study are also regularly exposed to violence or its imminent risk. In Maré, three criminal groups—two of which exploit the illegal drug market, and another is mainly an extortion racket—hold the monopoly on violence. Criminal groups were also present in the other two favelas during our fieldwork, and police was intermittent.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 workdays, 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 carries a stigma, i.e., a characteristic related to (unfair) generalizations, such as a negative stereotype, which might bear a grain of truth (Bordalo et al.,

⁵See 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.

2016) or that simply leads to antipathy in some people. If employer discrimination corresponds to the differential treatment dispensed to those who are recognized as belonging to a stigmatized group, we conceptualize *expected* discrimination as jobseekers' perceptions of the extent of that 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, or believe that outsiders simply dislike them, so they can expect discrimination. We consider similar jobseekers who live immediately outside favelas to be the nonstigmatized reference group. Relative to these nonfavela jobseekers, favela jobseekers might expect to receive lower callback rates and wage offers or to be treated with contempt. Hence, expecting discrimination can mean expecting lower returns to effort and more stress—the latter stemming from expecting hostile treatment, economic hardship, or feelings of injustice.

To illustrate how expected discrimination affects job application decisions more precisely, we introduce a simplified decision problem in which a jobseeker who lives in a favela decides whether and how to pursue a single job. We conceptualize expected discrimination as a belief distribution over employer types $d \in \mathbb{R}$, with $d > 0$ standing for discriminatory employers. Hence, the jobseeker is uncertain about how discriminatory the employer is. At a certain stage of the job application t , the jobseeker can pick a (costly) action $\mathbf{a} \in \mathcal{A}_t$ which can determine both task performance (broadly defined) and the perceived probability that the employer identifies the jobseeker as a favela resident at stage t , $\pi(\mathbf{a})$.⁶ If the jobseeker is perceived as a favela resident at stage t , she expects a continuation value $V(\mathbf{a}, d)$ with probability $p(\mathbf{a}, d)$. If not, she expects $V(\mathbf{a}, 0)$ with probability $p(\mathbf{a}, 0)$, i.e., she expects no discrimination at that stage. The value function V could also encode expected discrimination in stages after t , in case address becomes visible later, but we keep the notation simple for the moment. Then, a jobseeker at stage t might solve

$$\max_{\mathbf{a} \in \mathcal{A}_t} \mathbb{E}[\pi(\mathbf{a}) (p(\mathbf{a}, d)V(\mathbf{a}, d) - c(\mathbf{a}, d)) + (1 - \pi(\mathbf{a})) (p(\mathbf{a}, 0)V(\mathbf{a}, 0) - c(\mathbf{a}, 0))] \quad (1)$$

where $c(\mathbf{a}, d)$ is the cost of taking action $\mathbf{a}$. We let p and V decrease with d , so employers that are more discriminatory are less likely to advance candidates perceived to be from favelas and jobseekers perceive jobs with discriminatory employers to be worse. We may also let $c(\mathbf{a}, d)$ increase with d , so it may be more costly to perform well when one expects discrimination, e.g., due to psychological costs of controlling one's actions during a job interview when expecting

⁶For instance, higher effort may be read by an evaluator as both a signal of performance and a signal of not belonging to the stigmatized group.

discrimination.

This illustrates that expected discrimination affects job-seeking behaviors in proportion to (i) the probability the candidate assigns to being perceived as a favela resident π and (ii) the extent to which beliefs about d affect the continuation probabilities, benefits, and costs of pursuing a job. In our experiments randomizing expected address visibility, we manipulate π , varying the probability that the candidate would expect to suffer discrimination at that moment. In our Information Experiment, we manipulate expectations about continuation probabilities and the distribution of d . In Appendix E, we provide framework extensions that encode the elements present in each of our experiments.

Stereotype threat. Stereotype threat, the idea that being at risk of confirming some stereotype might impair performance, can enter the utility function in different ways depending on whose confirmation of the stereotype matters. If the jobseeker worries about confirming a stereotype (the stereotypical action a^t) *to themselves*, regardless of how observers perceive or treat them, we can model stereotype threat as an additional cost $c^t(a - a^t)$ appearing outside the expected value operator in Equation 1. We do not expect this type of stereotype threat, known as *self-concept threat* (Inzlicht and Schmader, 2012), to vary substantially across our treatment arms. First, because our treatments do not shift perceptions about what is stereotypical of a favela resident or about how diagnostic the candidate’s behavior is of their stereotype. Second, because we carefully control stereotype salience through our design, by mentioning address at least once to all applicants in each experiment. Alternatively, stereotype threat can stem from the fear of confirming a negative stereotype *in the minds of others*, known as *own reputation threat*. In our framework, that type of stereotype threat is a component of expected discrimination, which can be incorporated in the term $c(\mathbf{a}, d)$. Hence, our concept of expected discrimination partially overlaps with stereotype threat, but it is a much broader concept. For instance, an individual can be entirely confident in their abilities and not worry at all about confirming stereotypes, but still change their application effort or experience stress due to expected discrimination.

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. We designed the audit to be as relevant as possible to jobseekers in our supply-side experiments, whom we ask to predict its results.⁷ To ensure relevance, we only included jobs in regions within two hours by public

⁷Prior to our study, there was limited quantitative evidence on antifavela discrimination. Zanon et al. (2023) found substantial antifavela discrimination in Argentina using the incentivized résumé rating method (Kessler et al.,

transit; most jobseekers in our supply-side experiment with job experiences have worked in these same neighborhoods. A large share of jobseekers are also interested in, or have experience with, the sales positions included in the audit study because (i) about 40% applied to the sales jobs in our experiment and (ii) the positions included in the audit study were similar to the *last* formal job positions held by about one-third of the jobseekers. We summarize the study’s design and findings below. Appendix C provides additional details on the audit study design and implementation (including sample résumés) and presents descriptive evidence that the audited jobs overlap with our jobseeker sample in both geography and sector.

We created four fictitious worker profiles with completed high school, two male and two female. 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 each profile, we created two versions that differed only in name, email, phone number, and address—one from Maré and one from Bonsucesso, a nonfavela neighborhood adjacent to Maré. We picked random common names (among the 50 most popular by gender among adult workers), which are not particularly distinctive in terms of race, socioeconomic status (SES), or other dimensions subject to stereotypes.⁸

We applied to each job with two different profiles, randomizing whether we picked the version of the profile with the Maré or Bonsucesso address.⁹ We selected addresses that unambiguously mapped to either Maré or Bonsucesso while keeping the estimated commute time to any job constant. Maré is a widely recognized favela in Rio, so employers can immediately distinguish the neighborhood as a favela.

We collected job postings 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, and our results are similar if we recode the 27 requests for more information as callbacks.

2019). In Rio, an audit study (Westphal, 2014) found no antifavela discrimination, albeit with some geographical heterogeneity. 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.

⁸Only one of our names has a potential marker of low-SES, but our results do not change if we adjust for that. We discuss this in Appendix C.2, and Figure C.1 presents callback rates by name.

⁹Information 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.1).

We do not find evidence of statistically significant 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 equality of means regardless of whether we add controls or job fixed effects (Table C.2). We also see no evidence that antifavela discrimination varies with (i) whether jobs are located in richer or more distant neighborhoods or (ii) the type of hiring firm (e.g., insurance brokers, apparel, confidential).

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), that recruiters believe favela residents are *ceteris paribus* more likely to accept a job offer, offsetting callback differences caused by antifavela taste-based discrimination (Kessler et al., 2019), or that antifavela discrimination exists mainly in occupations excluded from our audit study (e.g., higher-status occupations in other sectors).¹⁰ 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. Finally, our experiment is underpowered to detect small discrimination rates: An ex post calculation for the minimum detectable effect puts it at about 3.6pp (22% of the nonfavela callback rate).

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 the callback rates we would find in our audit study. 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). Every jobseeker provided incentivized predictions for the callback rates we would find for Maré and Bonsucesso, and jobseekers from the other two favelas (Jacarezinho and Manguinhos) were also incentivized to predict the callback rates for their own favela and its adjacent nonfavela.¹¹ Our main measure of expected

¹⁰Although sales jobs represent a nontrivial share of formal employment in Rio (6.9% of all formal workers according to publicly available RAIS 2024 data), sales differ from other occupations in several ways. For instance, sales jobs are high-turnover, representing about 12% of all layoffs in 2024, and are not particularly specialized or high-status.

¹¹As our audit study only estimated callback rates for Maré and Bonsucesso, we incentivized these other predictions by initially stating that we knew the correct answer for only some of the questions, and that incentives would be paid based on those.

discrimination compares predictions about the jobseekers’ favela of residence with their adjacent nonfavela, but we reach similar conclusions if we instead focus on Maré-Bonsucesso comparisons (Figure A.2).

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). Jobseekers’ guesses for favelas are closer to the audit estimates but are still generally too optimistic: The average predicted callback rate for the jobseeker’s own favela is 30%—over 50% larger than the audit study estimates.

The bottom panel in Figure 2 shows the distribution of implied discrimination rates, i.e., the percent drop in 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.

Although we did not elicit jobseekers’ confidence level in predicting the audit study results, 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).¹² As it is reasonable to assume that predicted discrimination in callback rates correlates with expected discrimination in other moments (e.g., during an interview, or on the job), and as we do not have high-quality elicitation of these other differentials, we use predicted callback discrimination as our main proxy for baseline expected antifavela discrimination across our whole analysis.

While describing how people form beliefs about discrimination is beyond the scope of this paper, understanding the predictors of expected antifavela discrimination can help us interpret the experimental results. In our survey, we asked most jobseekers (69%) 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. Ordered from the most to the least agreement, the reasons are: 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

¹²Expected discrimination in the audit study also strongly correlates with two other expected discrimination proxies: a Likert measure of antifavela discrimination and a “personalized” discrimination measure comparing one’s own expected future employment probability against that of a similar jobseeker in the adjacent nonfavela (Figure A.3).

rich second-order beliefs about employers, expecting discrimination for both taste-based and statistical reasons. Further, those with higher reservation wages, who are not male, and who have a college degree tend to predict more antifavela discrimination (Table A.2).

2.5 Correlated Signals

Some characteristics—like race, skill level, and gender—are especially predictive of favela residence. We now discuss the role these characteristics may play in our experiments, motivating our heterogeneity analysis presented in Section 4.¹³

Race. Race and favela residence are strongly correlated, with 33% of favela residents self-identifying as white versus 57% of nonfavela residents, but race and address are generally perceived as two different sources of discrimination. Although jobseekers mention race as a reason why employers discriminate against favela residents, they also mention other factors related to reliability, safety, culture, skills, and dislike of favela residents (as discussed in Section 2.4). When we elicited expectations about racial and antifavela discrimination on a Likert scale, 84% of respondents believe firms discriminate “somewhat” or “a lot” against Black people, while 72% believe firms discriminate against favela residents to the same extent, indicating that expected racial discrimination may be stronger in general. Further, Table A.2 shows that mentioning racism as a main reason why firms do not hire from favelas is a strong predictor of expected antifavela discrimination, while respondent’s race is not, suggesting that jobseekers understand that employers may use address to proxy for race, regardless of their own race. These patterns suggest that jobseekers expect antifavela discrimination to operate beyond racial discrimination.

In our experiments, expected racial discrimination should not be a confounding factor, because in our design jobseekers always expect race to be visible—it needs to be declared in the application form and is visible in interviews. At the same time, the effects of expected address visibility may differ by race, due to common knowledge of the correlation between race and address. White jobseekers can “pass” as nonfavela residents when addresses are hidden, avoiding discrimination entirely, while nonwhite jobseekers may face discrimination regardless of address visibility. That is because, besides always expecting direct racial bias, they may also expect employers to infer favela residence from race. Hence, we could expect the effect of expected address visibility to be relatively smaller among nonwhite jobseekers.

Skill and Gender. Favela residence is negatively correlated with education: Favela res-

¹³In our context, first names are not typically perceived as distinctive in terms of race, but some names may signal lower SES. Therefore, some names may also signal favela residence. We discuss the role of names in the audit study in Section C.2 and also present heterogeneity analyses by whether a jobseeker’s name is distinctive of low-SES in Appendix A.

idents have approximately three fewer years of schooling than nonfavela residents, and such correlation is also common knowledge as 50% of jobseekers mention lower skill levels as a reason employers discriminate against favela residents.¹⁴ Additionally, males could in principle be subject to more antifavela discrimination since gang membership is a primarily male phenomenon (and 60% of jobseekers mention fear of violence as a reason why employers discriminate). Hence, when addresses are hidden, high-skilled jobseekers and women may more easily pass as nonfavela residents (though education will not be readily visible during interviews).

3 Experiment Design

This section presents our experimental design to test whether expected discrimination affects job-seeking behaviors. We first discuss our door-to-door recruitment procedures and sampling. Second, we present the experimental details on the Address Omission Experiment, which randomizes expected address visibility at the job application stage. Next, we discuss the Information Experiment, which provides market-level information about discrimination. Finally, we present the Interview Experiment, which randomizes expected address visibility at the interview stage. The experimental manipulations in the Address Omission Experiment and Interview Experiment shift expected address visibility in relation to the status quo, estimating the effects of “blinding” on job application rates and interview performance. The Information Experiment tests whether telling jobseekers that our audit study shows little evidence of antifavela discrimination reduces expected discrimination and encourages job applications. Figure 1 provides an overview of our experimental design. We preregistered these experiments and heterogeneity analysis across four dimensions (expected discrimination, race, skill, gender) before any randomization.¹⁵

3.1 Sample, Recruitment and Survey

Sample. Between March and June 2023, we conducted a door-to-door survey in three Rio de Janeiro favelas: Maré, Jacarezinho, and Manguinhos. We targeted individuals aged 18-40 who had completed high school (or were in their final year) and were actively seeking full-time formal sector employment. To minimize spillover effects and protect participant privacy, we enforced two sampling restrictions. First, we interviewed at most one person per household,

¹⁴See [IPS News](#) (09-2010).

¹⁵Appendix B discusses all deviations from preregistration and provides estimates of treatment effects on application effort and other secondary outcomes which were preregistered but are not discussed in the main text, for brevity.

and never interviewed multiple individuals simultaneously. Second, we excluded neighbors and family members of previous respondents based on address and last names. These restrictions ensure that treatment effects are not confounded by information sharing between participants. Our final sample consists of 2,167 jobseekers.¹⁶

Table A.3 presents summary statistics for our experimental sample. The average participant is 26 years old, 30% are men, and 22% self-identify as white. One-quarter had never held a job before the survey, while 32% reported currently working full- or part-time, predominantly in the informal sector. By design, our sample is younger and more educated than the general favela population. It is also somewhat less white than the average across all favela residents.

Surveyors. We hired local surveyors from each favela, because local residents can more easily establish trust and navigate the relevant social dynamics. Surveyors wore uniforms and badges, and introduced themselves as conducting an academic “labor market” survey for the University of British Columbia (UBC) research team—independent of local NGOs. While local NGOs helped us recruit and train surveyors and facilitated access to the territories, we chose not to associate the survey with any local NGO. That is because NGOs primarily provide social services (food assistance, job placement, community programs) in these favelas and we wanted to safeguard the purely academic nature of the research, without creating expectations of direct assistance.¹⁷

Survey. The survey had four main blocks. The first block asked about demographics and labor market experiences, without mentioning discrimination of any kind. The second block asked for permission to share the information from the first block with our partner HR firm (see details below). The third block asked about skills and certifications, including a one-minute basic math test. The surveyors then told participants that no information provided from the third block onward would be shared with the HR firm. Finally, the fourth block elicited incentivized beliefs about expected discrimination (where participants estimated callback rates in our audit study), followed by expected job market prospects, jobseekers’ perceptions of why employers would discriminate against favela residents, and contact information collection. We offered a participation incentive of R\$5 ($\approx$ 1 USD) and a chance to win another R\$500 (see Figure D.1 for photos of in-progress interviews). Appendix D.1.1 shows how we introduce the HR firm and elicit the callback rates predictions.

HR firm. We ran the HR firm ourselves. We emphasized a “partnership” between the

¹⁶Surveyors completed 2,392 valid surveys. Of these, 167 respondents declined to share their information with the HR firm, and an additional 58 provided invalid contact information, yielding 2,167 eligible participants.

¹⁷Large-scale surveys in favelas are rare and typically conducted by local NGOs. The most comparable data collection effort in Maré was a comprehensive census conducted before 2013. Manguinhos and Jacarezinho have weaker NGO networks and less frequent research activities. Given this context, surveyors emphasized the study’s academic nature and independence when talking to participants.

research team and the HR firm, which, while technically accurate, obscured the fact that we operated the HR firm ourselves. Our choice not to present the HR firm as part of the study was a form of limited deception to the extent that the jobseekers could not have anticipated that we would observe their interactions with the firm.¹⁸ Presenting the HR firm as separate from the research team helped keep participants' expectations about discrimination closer to those they would hold when applying to other formal jobs. Importantly, the survey team and the HR firm operated independently. Surveyors worked solely on door-to-door surveys and were not made aware that the research team operated the HR firm. Similarly, the HR firm staff had no connection with the survey team. This separation ensured that the HR firm's staff remained blind to treatment assignments in the pre-interview experiments. Finally, surveyors only mentioned the HR firm after the screening section, ensuring that the job referrals would not lead to selection into survey participation. In summary, the HR firm acted as an intermediary between our partner company and the jobseekers.

Hiring partner. 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.¹⁹ The partner firm committed to giving full consideration to and fast-tracking promising applicants recruited through our pipeline.

3.2 Supply-Side Experiments

Application pipeline. A few days after the door-to-door survey, the HR firm sent a WhatsApp message inviting all 2,167 participants to apply for the sales jobs. These participants took part in the Address Omission Experiment or the Information Experiment, as indicated in Figure 1. We invited all N=937 jobseekers who completed the application form to the interview. Of those, N=422 attended the interview and participated in the Interview Experiment. Appendix B provides more details on the rollout of each experiment.

Comparison to natural job search environments. We asked jobseekers whether they would like to share their data with a potential recruiter during the door-to-door survey. This approach resembles common recruiting practices: In our audit study, several HR firms contacted our profiles for jobs we had not directly applied, presumably based on data we entered

¹⁸Debriefing 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.

¹⁹These 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 these jobs.

on different websites. Beyond initial recruitment, all subsequent interactions followed standard practices in the Brazilian labor market: the HR firm contacted participants via WhatsApp;²⁰ the job application form mirrors those used by major Brazilian job platforms; and the interview process closely followed standard practices for entry-level sales positions. Hence, door-to-door recruitment provided a sampling frame but did not fundamentally alter jobseeker experience once invited to apply. Nevertheless, our sample may include some jobseekers who would not have applied to our sales jobs in a natural setting, for instance, because they are searching for openings in different occupations (which could attenuate our treatment effects, as these jobseekers would not apply regardless of treatment assignment). Additionally, the “invitation” framing may have increased participants’ perceived probability of success relative to cold applications.

3.2.1 Address Omission Experiment (N=1,303)

We designed the Address Omission Experiment to estimate the effects of expected discrimination on jobseekers’ behaviors by randomizing expected address visibility at the application stage.

Treatment. We randomly allocated N=1,303 jobseekers into one of three experimental conditions: (i) *Address Omission*, (ii) *Status Quo*, and (iii) *Known Address*. In the *Address Omission* arm, the HR firm’s invite stated that a home address “is not” needed (for applying). In the *Status Quo* and *Known Address* arms, the message instead stated that a home address “is also” needed (see Appendix D.2). The difference between *Status Quo* and *Known Address* is that, in *Status Quo*, the jobseeker filled in the address in the application form—the usual practice in our context (see Figure D.2). This allows us to observe how often applicants obfuscate their real addresses. In contrast, in *Known Address*, their favela neighborhood was already pre-filled, preventing obfuscation. Hence, comparing *Known Address* with *Address Omission* allows us to test how applicants’ behaviors change depending on whether they expect their favela address to be fully visible or not.

The HR firm sent the invitations in batches, up to ten days after participation in the door-to-door survey. We stratified the randomization by baseline expected discrimination, within each batch.

Application form. The application form started with a brief description of the available positions. Next, it asked applicants to confirm their name, phone, and address information (except in *Address Omission*, see Appendix D.3). It then proceeded as a standard application form,

²⁰82% of Brazilians use WhatsApp to communicate with companies [Opinion Box \(2025\)](#). Major employers have publicly stated using WhatsApp in their hiring processes (see [G1 \(12-13-2019\)](#)).

asking about demographics (date of birth and race)²¹, education, work experiences, courses and skills. Finally, it presented a full description of the three full-time jobs available: (i) (in-store) sales consultant, (ii) direct sales promoter, and (iii) direct sales supervisor (see Figure D.3 for the job descriptions).

Outcomes. Our main outcomes are whether the jobseeker completed the online application form and attended the job interview. 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 their address if their declared neighborhood is neither a favela nor the postal service neighborhood of the jobseeker's real address (surveyors recorded the latter after each survey).

3.2.2 Information Experiment (N=690)

The Information Experiment estimates the effects of expected discrimination on jobseekers' behaviors by providing information about the market-level discrimination measured in our audit study. Revealing the audit study findings has two potential advantages over our designs lowering expected address visibility. First, it provides cleaner experimental variation because it forgoes randomization in the application procedure, sidestepping confounders related to how jobseekers perceive the HR firm. Second, it is likely more policy-relevant since 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.

Treatment. We randomized N=690 participants into three treatment arms, embedded in the door-to-door survey: (i) *No Info*, in which we did not reveal any information, (ii) *Favela Info*, in which we informed participants about the favela callback rate (19.3%) after predicting the audit results, and (iii) *Full Info*, in which we informed participants about the favela and nonfavela callback rates (19.3% and 19.6%), revealing that we had found basically no discrimination in callback rates.²² As *Full Info* revealed not only the discrimination rate but also the favela callback level, we included *Favela Info* as an alternative control condition, so that knowledge of the favela callback level is constant across those two arms. See Figure D.4 for the graphs that surveyors used to convey the treatment. The HR firm invited survey respondents to apply for our partner's jobs using the *Status Quo* application form, which asked candidates to provide

²¹ Asking for race in job applications is common practice. In our audit study several application forms asked so.

²² We introduced the Information Experiment while we phased out the Address Omission Experiment. During the phase-out, there was an overlap of 183 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. Versions of the main tables and figures including the overlapping participants, which are very similar, can be found in Appendix B.

their home address.

Outcomes. In addition to the application and interview attendance outcomes, we also estimate the effects of information on immediate belief updates and address obfuscation. For the belief updates, we incentivized predictions of our HR firm’s callback rates for favela and non-favela résumés.²³ In addition, we sent participants a short endline survey two weeks after the experiment including Likert-scale questions about expected discrimination and the number of job applications submitted over the previous two weeks (see notes for tables A.4 B.2 for details).

3.2.3 Interview Experiment (N=422)

The Interview Experiment randomizes expected address visibility to test whether expected discrimination impacts interview performance.

Setup. We invited all jobseekers who completed the application form (N=937) for the interview, and all the N=422 candidates who attended the interview participated in the Interview Experiment. We conducted the interviews in a coworking space in downtown Rio de Janeiro. Attendees received a R\$25 ($\approx$ 5 USD) transport subsidy to cover travel costs. Upon arrival at the interview office, a receptionist greeted candidates and asked them to state their name, date of birth, and address, regardless of pre-interview experimental condition (so the jobseeker could obfuscate their address again; see Appendix D.5 for the procedure’s full description). To minimize the possibility of candidates interacting and discussing the experimental manipulations, we (i) scheduled interviews with 20-30 minutes gaps between start times, and (ii) only allowed a candidate to go up to the interview office if the preceding candidate had already left.

Treatment. Immediately before each interview, the receptionist told each candidate that their interviewer “will only know your name” (*Name-Only* condition) or “[...] your name *and* address” (*Name-and-Address* condition).²⁴ We provided interviewers with schedules containing only candidates’ names, never revealing specific addresses. Interviewers remained completely blind to the experimental manipulations throughout the study. We debriefed them only at the end of the experiment, ensuring that interviewer behavior did not confound our results.²⁵ Therefore, any treatment effects must arise from candidates’ beliefs about what the interviewer

²³There 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, while the surveyor stated that we knew the answers to only *some* of the questions.

²⁴Our design controls for an address “priming” mechanism because all candidates state their address to the receptionist before treatment assignment. Except for the possibility that this priming is dose-dependent, our experiment should isolate the effect of expecting the interviewer to know the candidate’s address.

²⁵We told candidates that interviewers “will know” their addresses, which became literally true by the study’s conclusion, once we debriefed the interviewers about the study’s purposes.

knew, not from interviewer discrimination. We instructed interviewers to maintain consistent behavior across all interviews and to adhere strictly to the script, and while interviewers may have inadvertently responded to candidate behaviors affected by the treatment (e.g., nervousness), any such responses must be first triggered by the candidate’s expectations.

We stratified the randomization by baseline expected discrimination within each batch—slower-than-expected candidate flow made more complex stratification schemes infeasible without severely compromising the logistics. Although we did not stratify by experimental condition in the previous experiments (Address Omission and Information Experiment), we show that Interview Experiment assignment is *ex post* balanced with respect to the upstream experimental assignments (Address Omission and Information Experiment conditions; see the first rows of Table A.9).

Selection into interview. Table A.5 compares baseline characteristics across participants in the three experiments. Relative to the Address Omission and Information Experiment samples, those in the Interview Experiment are younger, less likely to be employed, less likely to have prior formal sector experience, and more likely to report actively searching for jobs. This suggests that interview attendance was higher among jobseekers with more immediate employment needs, which would be expected in any job application procedure.

Interview and interviewers. Interviews lasted approximately 15-20 minutes and followed a structured script covering skills and fit for the position (see Appendix D.6 for the full script). We hired an experienced HR consultant to help us design our interview script and train our interviewers. We employed two female interviewers in their 30s. Regarding their race, we believe that jobseekers would perceive both interviewers as white. We hired the interviewers to be part of the HR firm, and they were completely disconnected from our team of surveyors. We trained them, emphasizing the following: (i) treating all candidates equally and with respect, (ii) adhering strictly to the script without asking personal questions or deviating into topics like candidates’ journey to the office, and (iii) maintaining consistent verbal and non-verbal behavior across all interviews, including tone of voice, eye contact, body language, and facial expressions. The interview script also reinforced these points.

Outcomes. Our main outcome is an index of self- and interviewer-assessed performance. Candidates filled the self-assessment in the reception area, after their interviews (see Appendix D.7). The receptionist stated that their responses would not affect their evaluation. Candidates indicated on 0-to-10 scales, (i) how well they performed overall, (ii) how nervous they were (reverse-coded as calmness), and (iii) how professionally they behaved. They answered these questions privately on a tablet, mitigating reporting bias (Bursztyn et al., *forthcoming*). The interviewers responded to these same questions once the interview was over (see Appendix

D.6). 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 the self- and interviewer-assessments (Anderson, 2008). As our primary aggregate performance measure, we average the self- and interviewer-assessment indexes, allowing us to extract a more accurate signal. Nearly all pairs combining two of the six components of our performance indexes are strongly correlated (Table A.6), indicating that they might indeed be all (noisy) indicators of performance.

3.3 Balance and Estimation

The realized treatment assignments generated comparison groups balanced across almost all pretreatment covariates, within what would be expected from truly random assignment (Tables A.7, A.8, and A.9).²⁶

To plot the average outcomes and test differences, we estimate a saturated model:

$$y_i = \alpha + \sum_{j \in T} \left[\beta^j t_i^j + \gamma^j t_i^j X_i \right] + \mu X_i + \varepsilon_i \quad (2)$$

where $y_i \in \{0, 100\}$ (to yield percentages) is, for instance, an interview attendance indicator, T is either $\{Status\ Quo, Address\ Omission, Known\ Address\}$ or $\{No\ Info, Favela\ Info, Full\ Info\}$, and t_i^j is a dummy for assignment to arm j . X_i is a vector of demeaned controls including covariates for which there is a potentially meaningful imbalance across treatment arms in an experiment, but the results are similar regardless of which controls are included (see Appendix A).²⁷ Thus, β^j is the covariate-corrected outcome level for outcome y in treatment arm j . In these regressions, we always use the sample of all individuals who were invited to apply, even when the outcome is downstream of the application decision, guaranteeing that selection is not a concern. For the Interview Experiment, as there are no imbalances at 5% and the outcomes are normalized, we simply estimate differences in means. We estimate robust standard errors.

To estimate whether effects depend on baseline characteristics, we define a heterogeneity cut for each variable: white *vs.* nonwhite, below *vs.* at-or-above median expected discrimination,

²⁶We expect a small number of random imbalances since we randomized 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. For the Interview Experiment, we also test for imbalances in the conditions in the two pre-interview experiments (Table A.9).

²⁷The threshold for inclusion is having a test for difference in average across any two treatment arms in an experiment with $p < 0.05$. X_i includes gender and skill in the Address Omission Experiment; residing in Manguinhos and age in the Information Experiment. 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).

male or not, or below vs. above median in our index of skill.²⁸ Then, we estimate:

$$y_i = \alpha + \sum_{j \in \tilde{T}} \left[\tilde{\beta}_1^j t_i^j h_i + \tilde{\beta}_2^j t_i^j (1 - h_i) \right] + \eta h_i + \mu X_i + \varepsilon_i \quad (3)$$

where $\tilde{T}$ are the same sets of treatments as T but omitting *Address Omission* (the lowest expected address visibility condition) or *No Info* and h_i is an indicator for being on a certain side of the heterogeneity cut. X_i is the same as in (2). $\tilde{\beta}_1^j$ then stands for the effect of j on those with $h_i = 1$ in relation to *Address Omission* or *No Info*, and $\tilde{\beta}_2^j$ stands for the effect on those with $h_i = 0$. In the main text, we focus on estimating heterogeneous effects on interview attendance since it is the “harder” application outcome and it determines selection into the Interview Experiment (further discussed in Section 4.3), but we present heterogeneity breakdowns for all main outcomes in Figures A.4 to A.7.²⁹

4 Results

4.1 Address Omission Experiment

On average, expected address visibility at the job application stage does not affect application decisions or interview attendance rates (Figure 3). If lower expected address visibility had led to lower exposure to expected discrimination, *Address Omission* should have the highest application rates, and *Known Address* the lowest, with *Status Quo* in between. Instead, we do not find any statistically significant differences in application or interview attendance rates. Application rates are 42% in *Status Quo* and *Address Omission*, and 45% in *Known Address*. Similarly, interview attendance only varies slightly: 19% in *Status Quo*, 20% in *Address Omission* and 17% in *Known Address*. These null average effects on interview attendance also indicate that manipulating expected address visibility does not lead candidates from any of the three treatment arms to be overrepresented in the Interview Experiment.

Failure to manipulate expected address visibility is unlikely to explain the null effects. In

²⁸We build the skill index as the simple average of the z-scores of: (i) the math test score, (ii) having completed high school, (iii) having some college, and (iv) surveyor-assessed communication skills.

²⁹We also estimate Equation (3) with additional interacted baseline covariates, reaching similar conclusions (Tables A.11 to A.13). Further, as our research design includes three experiments with multiple treatments and outcomes, testing for heterogeneity along these dimensions in each case separately can lead to false positives due to multiple hypothesis testing. To address this concern, we also present estimates of effect heterogeneity on key outcomes from a regression in which we stack data from all three experiments. We then present F-tests for whether we can reject effect homogeneity across all experiments at once, and the associated p -values can be Bonferroni-corrected by multiplying them by four, the number of heterogeneities tested (Table A.10). This analysis suggests that the heterogeneity by race is robust across experiments.

the *Status Quo* arm, when participants were free to obfuscate their addresses, 44% of those who finished their application forms (corresponding to 24% of those who were invited to apply) did so, suggesting that jobseekers indeed expected that the HR firm could discriminate based on address.³⁰ The *Known Address* treatment was also effective in increasing expected address visibility by preventing obfuscation since only one out of the 437 jobseekers in that condition provided a corrected address with an obfuscated neighborhood. In Section 5.1, we discuss mechanisms that could more effectively explain these results.

Heterogeneity. Table 2 presents treatment effects for each heterogeneity cut interacted with *Status Quo* and *Known Address* (Equation (3)). *Address Omission*, the condition with the lowest expected address visibility, is the omitted category. In column (1), we find evidence that the effects are heterogeneous by race. In *Address Omission*, white jobseekers are 10.9pp more likely ($p=0.02$) to attend the interview than nonwhite jobseekers. As address visibility increases, the racial attendance gap reverses. Among white jobseekers, *Status Quo* decreases attendance rates by 6.3pp ($p = 0.29$) and *Known Address* by 14.9pp ($p<0.01$), while these effects are close to zero among nonwhite jobseekers. We can reject that the effect of *Known Address* is the same among white and nonwhite jobseekers ($p=0.02$, last row of Table 2). This pattern suggests that hiding address information primarily benefits white jobseekers, who can more easily pass as nonfavela residents when their address is unknown. Nonwhite jobseekers may expect discrimination regardless of address visibility, either because race itself triggers discrimination or because employers may infer favela residence from race.

We find no statistically significant evidence for heterogeneous effects by baseline expected discrimination, gender, or skill. For those expecting high discrimination, who are not male, and who are more skilled, we estimate imprecise negative effects of *Known Address* relative to *Address Omission* of -3.6pp ($p=0.28$, column 2), -5.2pp ($p=0.11$, column 3), and -5.9pp ($p=0.11$, column 4), respectively. Point estimates of the effects of *Known Address* on the complementary groups are either smaller in magnitude or positive. These estimates are consistent with it being easier for women and the more skilled to “pass” as nonfavela residents when addresses are not visible (see Section 2.5), but we are underpowered to reject the nulls of no effect and of no effect heterogeneity.³¹

³⁰These rates refer to the sample of jobseekers with *Status Quo* application forms in the *Address Omission* Experiment. Table A.14 shows that expected discrimination is the main predictor of obfuscation across our three experiments, further suggesting that obfuscation is a strategic response to expected discrimination.

³¹Table A.15 further shows that there is no evidence of heterogeneous effects by names that signal low versus high-SES names.

4.2 Information Experiment

We now examine how providing information about market-level discrimination affects beliefs and application behaviors. The information treatments reveal callback rates from our audit study: *Favela Info* shows that the favela callback rate is 19.3%, while *Full Info* additionally reveals that the nonfavela callback rate is similar at 19.6%. Since jobseekers predict an average favela callback rate of 30%, *Favela Info* typically delivers a “negative” shock to posterior beliefs about callback rates. Even so, such a shock can have a nonmonotonic effect on application rates (as illustrated in Appendix E). *Full Info* adds a “positive” shock on top of *Favela Info*: Because most jobseekers overestimate discrimination, learning that the nonfavela callback rate is similar to the favela rate generally lowers expected antifavela discrimination, which should encourage more applications. Belief updating may be more complex, however, because jobseekers likely believe that the favela and nonfavela callback rates are correlated—both reflect conditions in the same labor market—so a signal about one rate can spill over into beliefs about the other. Under this “cross-updating,” *Favela Info* may lead jobseekers to revise both beliefs downward, and *Full Info*, by additionally revealing the (relatively low) nonfavela rate, may push both posteriors down further. Since lower expected callback levels could decrease job applications, even *Full Info* may not lead to a net increase in applications.

Belief updating. Average predicted favela callback rates hover around 35% across all three arms. However, both *Favela Info* and *Full Info* significantly reduce expected nonfavela callback rates—from 59% in *No Info* to 51% in *Favela Info* and 45% in *Full Info* ($p \leq 0.01$, see Figure 4), which suggests that there is cross-updating.³² The cross-updating pattern becomes clearer when we examine heterogeneity by baseline beliefs about the favela callback rate. For instance, among jobseekers who initially overestimated the favela callback rate, learning the observed nonfavela rate in *Full Info* leads to a further increase in the posterior belief about the favela callback rate (Figure A.8).

Combining the posterior beliefs about callback levels into expected discrimination rates, we find that, because of the cross-update, *Favela Info* also reduces posterior beliefs about expected discrimination, though less than *Full Info*. Average posterior expected discrimination in *No Info*, *Favela Info*, and *Full Info* is 36%, 29%, and 17%, respectively, with these differences significant at the 5% or 1% level (Figure 4). Hence, the interpretation of *Favela Info* and *Full Info* is somewhat muddled. Nevertheless, if expected discrimination is the major determinant of job application behavior, we should expect more pronounced encouragement effects in *Full*

³²Since 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 to be positively correlated across firms.

Info than in *Favela Info*.

Effects on behaviors. Despite the belief shifts, we do not find statistically significant differences in application and interview attendance rates across the information conditions (Figure 4). Application rates are 40% in *No Info* and 45% in both *Favela Info* and *Full Info*, with no statistically significant differences between them. Interview attendance rates also seem unaffected, hovering around 20% and indicating that those in the different arms of this Information Experiment are similarly represented in the interview sample. Two weeks after treatment delivery, we also find no significant effects on the self-reported total number of job applications submitted (Table A.4, Column 3), though expected discrimination remains relatively lower in *Full Info* compared to *Favela Info* ($p=0.06$) and *No Info* ($p=0.14$) (Column 2). We also cannot reject the null hypothesis of no effect on obfuscation rates.³³

Heterogeneity. As in the Address Omission Experiment, treatment effects vary by race. Compared to *No Info*, we estimate that *Full Info* reduces interview attendance by 15.7pp ($p=0.05$) among white jobseekers while having basically no effect among nonwhite jobseekers (a non-significant increase of 2.3pp; see column 1 in Table 3). In contrast, *Favela Info* has no effect on interview attendance regardless of race. This suggests that learning that there was no discrimination discourages white jobseekers. Several mechanisms, such as white jobseekers shifting their search away from sales jobs or an increase in the perceived number of competitors for each job, could explain this discouragement effect, although we lack data to directly test these hypotheses.

There is also some suggestive evidence of a differential effect of information by skill (Table 3, column 4), but not by baseline expected discrimination level (column 2) or gender (column 3). In relation to *No Info*, *Favela Info* increases attendance by 7.8pp ($p=0.13$) among the higher skilled while decreasing attendance by 6.6pp ($p=0.24$) among the lower skilled, with $p=0.06$ for the difference. *Full Info* induces a similar, but nonsignificant gap (3.8 for high-skilled and -8.0 for low-skilled, $p=0.12$ for the difference). Because *Full Info* and *Favela Info* produce similar effects in this case, such effects are more likely driven by the information that is common to both, i.e., about callback rate levels, than by expected discrimination.

³³We see some evidence that *Full Info* reduces obfuscation among those who underestimate favela callback rates, which is consistent with jobseekers strategically giving up on obfuscation since their initial beliefs (about its returns) were too pessimistic ($p=0.01$; Figure A.8). More generally, to better leverage belief heterogeneity, we might assume that the only way our treatments affect behaviors is through beliefs about callback rates and estimate the effect of those beliefs with two-stage least squares. This approach leads to similar conclusions, the main difference being a marginally significant negative effect of the posterior expected favela callback rate on obfuscation rates (see Table A.16).

4.3 Interview Experiment

Even if expected discrimination does not affect average application rates, it could still affect interview performance since there are many differences between application decisions and behavior in interviews. For instance, during the interview, the jobseeker must quickly adjust in response to the interviewer, who directly judges performance.

Effects on interview performance. When jobseekers expect that the interviewer would only know their name, the average aggregate performance index rises by 0.13SD (Figure 5), with the self-assessed index increasing by 0.17SD ($p<0.01$) and the interviewer-assessed index increasing by 0.09SD ($p=0.28$). While this suggests that the effects on self-assessments are larger on average, we cannot reject that they are equal ($p=0.33$). Notably, our point-estimates of the effects of reducing expected address visibility are positive across all index components. Starting with the self-assessed components, *Name-Only* increases perceived overall performance by 0.13SD ($p=0.18$), calmness by 0.15SD ($p=0.13$), and professionalism by 0.2SD ($p=0.02$). Our estimates of effects on interviewer’s assessments are less precise and somewhat smaller: 0.06SD for overall performance ($p=0.5$), 0.12SD for calmness ($p=0.2$), and 0.07SD for professionalism ($p=0.5$).

Heterogeneity. For candidates expecting high discrimination at baseline, *Name-Only* increases performance by over 0.2SD on the aggregate, self- and interviewer-assessments ($p<0.01$, $p<0.01$, and $p=0.04$, respectively), whereas there is no evidence that *Name-Only* impacts performance among candidates expecting low discrimination (Table 4, columns 4-6). We also reject that the effects on aggregate and interviewer-assessed performance are the same regardless of baseline expected discrimination ($p=0.04$ and $p=0.05$, respectively). Hence, shutting down expected discrimination improves performance particularly among those who expect higher discrimination at baseline, even when interviewers do not have the information to discriminate. This pattern reveals that expected discrimination acts as a self-fulfilling prophecy: Believing one will face bias impairs performance, worsening outcomes independent of whether discrimination actually occurs.³⁴

Regarding heterogeneity by race, we find that *Name-Only* increases white candidates’ performance by over 0.3SD in the self- ($p=0.02$) and interviewer’s ($p=0.08$) assessments (Table 4, Columns 2-3). Among nonwhite candidates, *Name-Only* has a smaller effect on self-assessed performance (0.12SD, $p=0.09$) and no apparent effect on interviewer-assessed performance (0.02SD, $p=0.83$). We have only weak evidence that the effects on aggregate performance

³⁴Expected 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. Among candidates expecting high discrimination, only half as many score above the 90th percentile of interviewer-assessed performance in *Name-and-Address* compared to *Name-Only* ($p=0.09$; Figure A.9).

differ across racial lines ($p=0.09$), and we can't reject that the effects on each of the two performance components are the same regardless of race. This differential effect by race suggests that the same mechanism discussed in the context of the Address Omission Experiment is at play here. That is, telling white candidates that the interviewer would only know their name makes it easier for them to pass as a nonfavela resident, while nonwhite candidates might expect substantial discrimination regardless—because they may always expect racial discrimination and interviewers to stereotype them as favela residents.³⁵

We now discuss heterogeneity by skill and gender (see Table 4, columns 7 to 12). Among the more skilled, there is no evidence that *Name-Only* improves performance (0.06SD in the aggregate index, $p=0.39$). Among the less skilled, *Name-Only* increases aggregate performance by 0.19SD ($p=0.03$). Effect estimates are even more similar across gender lines—0.18SD ($p=0.12$) for males and 0.11SD ($p=0.1$) for nonmales. We cannot reject that the effects are the same across these heterogeneity cuts (last row of Table 4).³⁶

Selection into the interview and external validity. Since the interview is downstream of the application stage experiments (Address Omission and Information experiments), we may worry about whether the estimated average treatment effects of *Name-Only* differ from those that would have been revealed under more natural conditions. There are two mechanisms through which the existence of the application-stage treatments could affect the interpretation of the interview effects. First, in what we call a composition mechanism, the non-status quo experimental conditions may lead different types of individuals to self-select into attending the interviews. In fact, in the previous sections, we saw evidence that white jobseekers were (i) more likely to attend interviews under *Address Omission* than under *Known Address* and (ii) less likely to attend the interview under *Full Info* than under *No Info* condition. Second, in what we call a preconditioning mechanism, the non-status quo treatments themselves might change how people react to *Name-Only* (e.g., those under *Full Info* might have been less affected by *Name-Only* if *Full Info* had led them to believe interviewers were less likely to engage in antifavela discrimination). Of course, both mechanisms may interact if the preconditioning mechanism depends on baseline characteristics. Given the relatively small sample size in the Interview Experiment, separating the roles of the composition and preconditioning mecha-

³⁵ 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.

³⁶ We find no evidence of heterogeneous effects by names that signal low-SES (Table A.17). Additionally, Table A.18 shows that the choice to obfuscate address either at the application or the interview stages is not particularly predictive of the effects of *Name-Only* on interview performance.

nisms is impractical.³⁷ One specific composition concern is whether treatments that encourage white applications (e.g., Address Omission) filter in lower-quality white interviewees, potentially confounding race-specific interview effects. Small sample sizes preclude definitive tests of this mechanism, though the concentration of effects in our most externally valid subsample (discussed below) suggests composition effects do not primarily drive our results.

We approach the problem caused by the selection and preconditioning mechanisms by focusing on estimating whether the effects of *Name-Only* are the same in more vs. less “externally valid” subsamples (i.e., closer or further away from the status quo). If the effects are similar, it would mean that the composition and preconditioning mechanisms are not very important for determining the ATEs. If they are different, we should put more trust on the ATEs estimated for the more externally valid subsample when thinking about policy implications. 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.”³⁸

We find that the effects of *Name-Only* are either similar or larger in the more externally valid subsample. We estimate that the effect of *Name-Only* on aggregate performance is 0.23SD ($p < 0.01$) for the group with more external validity and 0.03SD ($p = 0.7$) for that with less—see Table 1. There is some evidence that these effects are statistically different ($p = 0.09$), but the difference is not in the direction that would threaten the external validity of our main findings. 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 effects on interviewer perceptions. *Name-Only* has a strong effect on interviewer evaluations for this group (0.35SD, $p < 0.01$) and a negative, nonsignificant effect among the subsample for which our findings are less externally valid. Hence, it appears that the composition and preconditioning effects might have led, in total, to a lower estimate of the average treatment effect of *Name-Only* on performance, due to smaller effects on interviewer assessments in the less externally valid subsamples.

³⁷For instance, even if we only take into account two baseline characteristics, such as race and skill, we would need to estimate the effects of *Name-Only* for different combinations of race and skill for each of the possible pre-interview experimental conditions, rapidly running into comparisons based on 20 observations or less.

³⁸This strategy generates an approximately even split of participants, making the test more powerful. In Table A.19, we estimate the effects of *Name-Only* for a finer partition of the pre-interview treatments, obtaining similar results.

5 Discussion

5.1 Mechanisms

Our main finding indicates that expected discrimination hurts interview performance, but does not change average application behaviors. In this subsection, we discuss potential mechanisms that could explain our main finding in light of our data and the broader literature.

5.1.1 Interview effects

In job interviews, performance is particularly sensitive to stress: stakes are high, time is limited, and behavior is immediately observed and judged face-to-face. Interviewees might “choke under pressure,” a well-documented phenomenon in which people perform below their demonstrated ability precisely when stakes are high (Baumeister, 1984; Böheim et al., 2019; Godlonton, 2020). Expected discrimination can further exacerbate stress, as documented in the psychology and human biology literatures (see Schmitt et al. (2014) and Berger and Saranyai (2015) for meta-studies, respectively). Consistent with this, we see that the effect of *Name-Only* on self-assessed calmness is 0.4SD larger among those expecting high discrimination than among those expecting lower discrimination (Figure A.7, $p=0.04$). Also consistent with choking under pressure is the fact that candidates self-perceive as behaving 0.2SD less professionally under *Name-and-Address*, suggesting that expecting their address to be visible disrupts their performance. Although the aforementioned patterns in candidate behavior are consistent with choking under pressure, we cannot completely rule out alternative explanations. For instance, it is possible that lowering expected discrimination through *Name-Only* causes the perceived returns to exerting effort in the interview to increase, leading to higher effort and better performance. In addition, expected discrimination concerns may just be much more salient when face-to-face. As we do not have measures of beliefs about returns to effort or of salience (both of which would be less natural in our setting), we cannot directly speak to such theories.

5.1.2 Null Effects on Application Decisions

Our null average effects on application decisions suggest that they are largely inelastic to expected discrimination at the application stage. One possible explanation for this inelasticity is that most jobseekers believe that they might suffer address discrimination anyway in later stages of the matching procedure. While we do not have the data to verify this mechanism in the Address Omission Experiment, it seems like a potential candidate for explaining the null effects in the Information Experiment. Table B.2 (column 3) shows that those who received *Full*

Info or *Favela Info* expect to suffer on-the-job address-based discrimination with similar odds. Another possible explanation for this inelasticity is that jobseekers use a heuristic of applying to all jobs that could be a “good fit.” Accordingly, survey questions that we introduced along with the Information Experiment suggest many jobseekers apply broadly and indiscriminately: 68% agree one should apply for all possible postings, and 84% agree one should not ruminate about employer discrimination. With 68% of our sample unemployed and 49% actively searching, the low cost of completing an online application may lead jobseekers to apply regardless of discrimination concerns. However, given our suggestive evidence that white jobseekers do respond to expected discrimination, the explanations above may not provide a full account of our findings.

We can also speak to some mechanisms explaining the null effects of providing market-level information in the Information Experiment, specifically. First, revealing market-level statistics may simply not be persuasive enough to shift behavior. A lack of confidence on posterior beliefs about callback rates suggests this might play a role in explaining the results: Only one-third of participants reported feeling “very” or “extremely” confident in their predictions about what callback rates the HR firm would implement. Second, we can rule out the possibility that expected discrimination leads people to apply to more jobs outside the ones in our study, as we find no evidence of increased applications overall (Table A.4, column 3). Still, people could have increased the type or quality of the jobs they apply for (paralleling the findings in Kiss et al. 2023 and Belot et al. 2019), or changed *how* they search for jobs, which we do not observe. Third, because (i) shifting expected callback levels can have nonmonotonic effects on search effort and (ii) favela jobseekers also use the nonfavela callback rate to update about the favela callback rate, the effects of lower expected discrimination driven by *Full Info* may be confounded and dampened by effects on callback levels. While we cannot offer empirical evidence about this mechanism, we illustrate it with a simple search framework in Appendix E.

5.2 Remarks and Policy Implications

The different effects across environments (application vs. interview) suggest that interventions targeting expected discrimination may be most effective when focused on high-stakes, in-person interactions rather than on earlier application stages. We emphasize, however, that our null application effects should not be interpreted as evidence that expected discrimination is never relevant for determining application behaviors. As 84% of our sample overestimates market-level discrimination, these beliefs may shape job search behaviors at margins we do not observe, such as which industries or firms jobseekers target. In addition, 42% of applicants obfuscate their true address, suggesting jobseekers do apply strategically to deal with expected

discrimination.

Our interview findings can also have further real-world implications. For instance, negative interview experiences can discourage future applications: Jobseekers who feel they performed poorly may become reticent to apply for other jobs requiring formal interviews. Moreover, our experimental design likely provides a lower bound on the effects of expected address visibility on performance, since interviewers were blind and trained to behave the same in all interviews, eliminating discriminatory interviewer reactions. Finally, even absent hiring consequences, expected discrimination imposes psychological costs (Pascoe and Smart Richman, 2009; Schmitt et al., 2014; Berger and Sarneyai, 2015), so the negative interview experiences we document may undermine jobseekers’ welfare regardless of employment outcomes.

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 *generally* encourage more applications. Since there is also evidence that blinding can backfire when it leads 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.

Nevertheless, 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 employers make jobseekers aware of credible blinding or antidiscrimination policies, they might extract a more positive (and potentially more accurate) ability signal when interviewing people who expect discrimination. Designing settings that do not allow for correlated sources of discrimination to inform each other (e.g., audio-only, text-only, or metaverse interviews) should help ensure effectiveness 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. (2024).

5.2.1 Generalizability

Our findings apply most directly to the population in our study, i.e., relatively young and educated favela jobseekers. The extent to which our policy implications generalize beyond our setting depends on whether the conditions and mechanisms behind our findings are likely to operate in other contexts and occupations, which we discuss next.

Information provision. Would the provision of information about actual market-level discrimination rates also fail to encourage more applications in other settings? Some limited evidence from other settings suggests that discrimination may be typically overestimated ([Haaland and Roth, 2023](#); [Aksoy et al., 2023](#); [Angeli and Lowe, 2024](#)), so that accurate information would in principle be actionable. Nevertheless, as discussed above, other studies of information provision also failed to find an increase in the number of jobs applied for (rather finding that people change which jobs they apply for, see [Kiss et al. 2023](#) and [Belot et al. 2019](#)) and we can only speculate about the reasons why information provision failed to encourage applications in our study. In a different setting, e.g., in which jobseekers hold tight priors and update more strongly based on market-level data, it is still possible for information to affect application behavior.

Expected discrimination and interview performance. If choking under pressure explains, at least partially, why expected discrimination hurts interview performance, we believe this effect should generalize to other contexts. First, because there is a body of research showing that expected discrimination appears to induce stress in several different settings and due to different kinds of discrimination (e.g., racial or against people with disabilities, see [Berger and Sarnyai 2015](#) and [Schmitt et al., 2014](#)). And second, because choking under pressure can happen across many different settings, from professional interviews ([Godlonton, 2020](#)) to world-level sport competitions ([Böheim et al., 2019](#); [Harb-Wu and Krumer, 2019](#); [Teeselink et al., 2020](#)), suggesting that it is a universal phenomenon.

Finally, the extent to which blinding procedures improve interview performance depends on whether blinding is both credible and effective in each setting. Blinding should be less effective when the stigmatized trait is correlated with other visible signals that remain available at the interview stage, or when the interview format necessarily reveals information that serves as a proxy for the concealed attribute.

6 Conclusion

Labor market discrimination is pervasive, yet we know far less about how jobseekers' expectations of discrimination shape their behavior. This paper shows that expected discrimination does not change application behaviors but negatively impacts performance at job interviews, operating as a self-fulfilling prophecy. We document a novel mechanism through which expected discrimination becomes self-fulfilling: By impairing performance in high-stakes interviews, rather than through reduced human capital investment ([Coate and Loury, 1993](#)) or interacting with on-the-job discrimination ([Glover et al., 2017](#)). Expected discrimination may therefore contribute

to persistent labor market inequalities through channels beyond those previously documented.

Our paper leaves several open avenues for future research. First, why does expected discrimination affect interview performance but not application behaviors? Experiments could vary interview pressure or format (face-to-face versus online) to isolate the mechanisms. Second, how do effects vary when people have multiple correlated stigmatized characteristics that differ in visibility? Our suggestive heterogeneity by race indicates that accounting for correlated characteristics with different visibility may be important, which speaks to the concept of intersectionality (Carvalho et al., 2022; Crenshaw, 1989). Third, to what extent people behave strategically as a response to expected discrimination in face-to-face interactions? While research documents strategic behaviors at “cold” stages (e.g. Charness et al. 2020), we have less evidence of verbal or non-verbal cues used during in person interactions, which relates to a large literature in psychology on coping (Folkman et al., 1986). Finally, could providing information about actual discrimination affect interview performance, even when we do not see effects on application decisions?

References

- Abebe, Girum, A Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, Simon Quinn, and Forhad Shilpi, “Matching frictions and distorted beliefs: Evidence from a job fair experiment,” *The Economic Journal*, 2025, 135 (671), 2089–2121.
- 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.
- Alston, Mackenzie, “The (perceived) cost of being female: An experimental investigation of strategic responses to discrimination,” *Working paper*, 2019.
- Amer, Abdelrahman, Ashley C Craig, and Clémentine Van Effenterre, “Decoding gender bias: The role of personal interaction,” Technical Report, IZA DP No. 17077 2024.
- 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 Vecchi**, “Does artificial intelligence help or hurt gender diversity? Evidence from two field experiments on recruitment in tech,” Working Paper 10996, CESifo 2024.
- 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,” *Journal of Labor Economics*, 2025, 43 (3), 000–000.
- 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 Gröbl, 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, Peter Hull, and Alex Imas**, “Systemic discrimination: Theory and measurement,” *Quarterly Journal of Economics*, 2025, 140 (3), 1743–1799.
- Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer**, “Stereotypes,” *Quarterly Journal of Economics*, 2016, 131 (4), 1753–1794.

- Boring, Anne, Katherine Coffman, Dylan Glover, and Maria José González-Fuentes**, “Discrimination, Rejection, and Willingness to Apply: Effects of Blind Hiring Processes,” Working Paper 2025. Unpublished manuscript.
- 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.
- Bursztyn, Leonardo, Ingar K. Haaland, Nicolas Röver, and Christopher Roth**, “The Social Desirability Atlas,” *Journal of Political Economy Microeconomics*, forthcoming.
- 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?,” *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.
- Del Carpio, Lucia and Thomas Fujiwara**, “Do Gender-Neutral Job Ads Promote Diversity? Experimental Evidence from Latin America’s Tech Sector,” Working Paper 2025. Unpublished manuscript.
- Delfino, Alexia**, “Breaking gender barriers: Experimental evidence on men in pink-collar jobs,” *American Economic Review*, 2024, 114 (6), 1816–1853.
- Dickerson, Andy, Anita Ratcliffe, Bertha Rohenkohl, and Nicolas Van de Sijpe**, “Anticipated labour market discrimination and educational achievement,” *Journal of Economic Behavior & Organization*, 2024, 222, 375–393.
- 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.

- Ferreira, Pedro Cavalcanti, Alexander Monge-Naranjo, and Luciene Torres de Mello Pereira**, “Of Cities and Slums,” *Journal of Political Economy*, 2025, 133 (9), 2693–2734.
- Flory, Jeffrey A, Andreas Leibbrandt, Christina Rott, and Olga B Stoddard**, “Signals from On High and the Power of Growth Mindset: A Natural Field Experiment in Attracting Minorities to High-Profile Position,” Working Paper 5807, Faculty Publications 2022.
- Folkman, Susan, Richard S Lazarus, Christine Dunkel-Schetter, Anita DeLongis, and Rand J Gruen**, “Dynamics of a stressful encounter: cognitive appraisal, coping, and encounter outcomes,” *Journal of Personality and Social Psychology*, 1986, 50 (5), 992.
- Glover, Dylan**, “Job search and intermediation under discrimination: Evidence from terrorist attacks in France,” Working Paper 164, Chaire Securitisation des Parcours Professionnels 2019.
- , **Amanda Pallais, and William Pariente**, “Discrimination as a self-fulfilling prophecy: Evidence from French grocery stores,” *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.
- Goldsmith, Arthur H, Stanley Sedo, William Darity Jr, and Darrick Hamilton**, “The labor supply consequences of perceptions of employer discrimination during search and on-the-job: Integrating neoclassical theory and cognitive dissonance,” *Journal of Economic Psychology*, 2004, 25 (1), 15–39.
- 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 Joseph E Stiglitz**, “Equilibrium fictions: A cognitive approach to societal rigidity,” *American Economic Review*, 2010, 100 (2), 141–146.

- **and Priyanka Pandey**, “Discrimination, social identity, and durable inequalities,” *American Economic Review*, 2006, 96 (2), 206–211.
- **and —**, “Making up people—The effect of identity on performance in a modernizing society,” *Journal of Development Economics*, 2014, 106, 118–131.
- Ibañez, Marcela and Gerhard Riener**, “Sorting through affirmative action: Three field experiments in Colombia,” *Journal of Labor Economics*, 2018, 36 (2), 437–478.
- Instituto Brasileiro de Geografia e Estatística**, “Censo Demográfico 2010: Resultados do Universo — Agregados por Setor Censitário,” Rio de Janeiro: IBGE 2011. Census-tract-level dataset, spreadsheets for the state of Rio de Janeiro (accessed June 2026).
- Inzlicht, Michael and Toni Schmader**, *Stereotype threat: Theory, process, and application*, Oxford University Press, 2012.
- 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,” *Quarterly Journal of Economics*, 2022, 137 (4), 1963–2036.
- Kuhn, Peter and Kailing Shen**, “What happens when employers can no longer discriminate in job ads?,” *American Economic Review*, 2023, 113 (4), 1013–1048.
- Lang, Kevin and Michael Manove**, “Education and labor market discrimination,” *American Economic Review*, 2011, 101 (4), 1467–1496.
- Leibbrandt, Andreas and John A List**, “Do equal employment opportunity statements encourage racial minorities? evidence from a large natural field experiment,” *European Economic Review*, 2025, 174, 104987.

- Lepage, Louis-Pierre, Xiaomeng Li, and Basit Zafar**, “Anticipated Gender Discrimination and Grade Disclosure,” Working Paper 30765, National Bureau of Economic Research 2022.
- , —, and —, “Anticipated Discrimination and Major Choice,” Working Paper 33680, National Bureau of Economic Research 2025.
- 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.
- Lundberg, Shelly J and Richard Startz**, “Private discrimination and social intervention in competitive labor market,” *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.
- Opinion Box**, “WhatsApp no Brasil 2025,” <https://content.app-us1.com/JY8yY/2025/07/01/3389a59f-4de8-465b-b040-0cec3fbadc6f.pdf> July 2025. Accessed on August 9, 2026.
- Pager, Devah and David S Pedulla**, “Race, self-selection, and the job search process,” *American Journal of Sociology*, 2015, 120 (4), 1005–1054.
- , **Bart Bonikowski, and Bruce Western**, “Discrimination in a low-wage labor market: A field experiment,” *American Sociological Review*, 2009, 74 (5), 777–799.
- Pascoe, Elizabeth A and Laura Smart Richman**, “Perceived discrimination and health: a meta-analytic review.,” *Psychological bulletin*, 2009, 135 (4), 531.
- Pethig, Florian and Julia Kroenung**, “Biased humans,(un) biased algorithms?,” *Journal of Business Ethics*, 2023, 183 (3), 637–652.

- 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.
- Riach, Peter A and Judith Rich**, “Field experiments of discrimination in the market place,” *The Economic Journal*, 2002, 112 (483), F480–F518.
- Rich, Judith**, “What do field experiments of discrimination in markets tell us? A meta analysis of studies conducted since 2000,” Technical Report 2014.
- Ruebeck, Hannah**, “Perceived discrimination at work,” 2024.
- Schmitt, Michael T, Nyla R Branscombe, Tom Postmes, and Amber Garcia**, “The consequences of perceived discrimination for psychological well-being: a meta-analytic review,” *Psychological bulletin*, 2014, 140 (4), 921.
- Scottini, Lucas Costa**, “O que o Nome nos ensina? Padrões sociais e raciais de nomes e sobrenomes e performance escolar em São Paulo.” PhD dissertation, Universidade de São Paulo 2011.
- Shukla, Soumitra**, “Making the Elite: Coded Discrimination at Top Firms,” 2024.
- 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.
- 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?,” *Working paper*, 2023.

Table 1: Interview Treatment Effects by Treatment Conditions Before Interview

	(1) Aggregate performance index	(2) Interviewer-assessed performance index	(3) Self-assessed performance index
<i>Name-Only</i> × <i>Status Quo</i> × <i>non-Full Info</i>	0.23*** (0.09)	0.35*** (0.12)	0.12 (0.10)
<i>Name-Only</i> × <i>All other pre-interview conditions</i>	0.03 (0.08)	-0.13 (0.11)	0.19** (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). *Name-Only* × *Status Quo* × *non-FullInfo* corresponds to *Name-Only* interacted with an indicator for participants who both followed a *Status Quo* application procedure and did not receive *Full Info* (i.e., pooling the *Status Quo* × *No Info* and *Status Quo* × *Favela Info*). The second coefficient corresponds to all other pre-interview paths. Regressions fully control for treatment assignment in previous experiments. The last row compares the two regression coefficients displayed in each column. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Treatment Effects Heterogeneity in the Address Omission Experiment

	(1) Show Up (%)	(2) Show Up (%)	(3) Show Up (%)	(4) Show Up (%)
<i>Status Quo</i> $\times h$	-6.34 (6.02)	0.40 (3.42)	0.42 (4.47)	-0.06 (3.98)
<i>Known Address</i> $\times h$	-14.88*** (5.71)	-3.58 (3.30)	2.16 (4.59)	-5.87 (3.71)
<i>Status Quo</i> $\times (1 - h)$	0.88 (3.01)	-3.49 (4.41)	-1.75 (3.38)	-1.23 (3.72)
<i>Known Address</i> $\times (1 - h)$	0.64 (2.96)	-1.81 (4.47)	-5.17 (3.27)	-0.18 (3.78)
<i>h</i>	10.93** (4.85)	1.31 (4.01)	-7.22* (3.84)	3.02 (4.65)
Observations	1303	1303	1303	1303
<i>h</i>	White	High $\mathbb{E}[disc.]$	Male	More skilled
<i>Address Omission</i> mean when $h = 0$	17.13	18.79	22.34	18.50
<i>Status Quo</i> = <i>Known Address</i> when $h = 1$, p-val	0.12	0.23	0.72	0.14
<i>Status Quo</i> = <i>Known Address</i> when $h = 0$, p-val	0.93	0.70	0.27	0.77
Effect of <i>Status Quo</i> constant across h , p-val	0.28	0.48	0.70	0.83
Effect of <i>Known Address</i> constant across h , p-val	0.02	0.75	0.19	0.28

Note: OLS estimates for the effects of *Status Quo* and *Known Address* on show-up rates for groups with $h_i = 1$ and $h_i = 0$. Regressions control for male and skill index. The scalars report p-values for tests of equality: whether treatment effects have equal effects within each group (rows 3-4), and whether treatment effects differ by h (rows 5-6). Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Treatment Effects Heterogeneity in the Information Experiment

	(1) Show Up (%)	(2) Show Up (%)	(3) Show Up (%)	(4) Show Up (%)
<i>Full Info</i> $\times h$	-15.66** (7.95)	-4.90 (5.79)	-3.37 (6.49)	3.84 (5.21)
<i>Favela Info</i> $\times h$	1.09 (8.97)	2.73 (5.88)	1.70 (6.63)	7.79 (5.15)
<i>Full Info</i> $\times (1 - h)$	2.34 (4.33)	1.09 (5.05)	-0.48 (4.66)	-7.96 (5.63)
<i>Favela Info</i> $\times (1 - h)$	1.33 (4.16)	0.12 (4.92)	1.41 (4.57)	-6.58 (5.57)
<i>h</i>	5.57 (7.23)	2.99 (5.52)	-4.34 (5.79)	-2.01 (5.60)
Observations	690	690	690	690
<i>h</i>	White	High $E[disc.]$	Male	More skilled
<i>No Info</i> mean when $h = 0$	18.60	18.33	21.43	21.59
<i>Full = Favela</i> when $h = 1$, p-val	0.03	0.18	0.42	0.47
<i>Full = Favela</i> when $h = 0$, p-val	0.81	0.85	0.68	0.78
Effect of <i>Full Info</i> constant across h , p-val	0.05	0.43	0.72	0.12
Effect of <i>Favela Info</i> constant across h , p-val	0.98	0.73	0.97	0.06


Note: OLS estimates for the effects of *Full Info* and *Favela Info* on show-up rates for groups with $h_i = 1$ and $h_i = 0$. Regressions control for mangueiros and age. The scalars report p-values for tests of equality: whether treatment effects have equal effects within each group (rows 3-4), and whether treatment effects differ by h (rows 5-6). Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Treatment Effects Heterogeneity in the Interview Experiment

	(1) Aggregate performance index	(2) Interviewer- assessed performance	(3) Self- assessed performance	(4) Aggregate performance index	(5) Interviewer- assessed performance	(6) Self- assessed performance	(7) Aggregate performance index	(8) Interviewer- assessed performance	(9) Self- assessed performance	(10) Aggregate performance index	(11) Interviewer- assessed performance	(12) Self- assessed performance
<i>Name-Only</i> $\times h$	0.31** (0.13)	0.31* (0.17)	0.32** (0.13)	0.22*** (0.08)	0.22** (0.10)	0.23*** (0.08)	0.18 (0.11)	0.10 (0.17)	0.25** (0.12)	0.06 (0.08)	0.03 (0.10)	0.10 (0.08)
<i>Name-Only</i> $\times (1 - h)$	0.07 (0.07)	0.02 (0.09)	0.12* (0.07)	-0.02 (0.09)	-0.11 (0.13)	0.07 (0.09)	0.11 (0.07)	0.08 (0.09)	0.14* (0.07)	0.19** (0.09)	0.12 (0.12)	0.25*** (0.10)
<i>h</i>	-0.05 (0.10)	0.06 (0.13)	-0.15 (0.11)	-0.06 (0.09)	-0.10 (0.12)	-0.03 (0.09)	0.09 (0.10)	0.08 (0.14)	0.11 (0.11)	0.32*** (0.08)	0.51*** (0.11)	0.12 (0.09)
Observations	422	422	422	422	422	422	422	422	422	422	422	422
<i>h</i>	White	White	White	High $E[disc.]$	High $E[disc.]$	High $E[disc.]$	Male	Male	Male	More skilled	More skilled	More skilled
<i>Name-and-Address</i> mean when $h = 0$	.01	-.01	.04	.04	.06	.02	-.02	-.02	-.03	-.17	-.28	-.06
Effect is constant across h , p-val	0.09	0.14	0.20	0.04	0.05	0.22	0.61	0.90	0.44	0.29	0.57	0.23


Note: OLS estimates for the effects of *Name-Only* on interview performance indexes for groups with $h_i = 1$ and $h_i = 0$. Sample restricted to those who showed up to the interview. The scalar reports the p-value for testing whether the treatment effect differs by h . Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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 Address Omission Experiment happened from March to May 2023, and the Information Experiment from June to July 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: Address Omission Experiment: No Differences in Application Rates Across Arms




Note: Bars show the covariate-adjusted share of all invited jobseekers in the Address Omission Experiment reaching each stage of the application process, in percent (β^j from Equation 2). *Applied* means finishing the online application form, and *Show Up* means showing up at the interview. Sample size in each arm is shown at the bottom of each bar. Vertical error bars display 95% confidence intervals based on robust standard errors, and brackets show p-values for pairwise tests of equality between arms.

Figure 4: Information Treatment Shifts Beliefs, But Not Interview Show-up



Note: Bars show the covariate-adjusted mean of each outcome by experimental arm, in percent (β^j from Equation 2). The *Beliefs* panel displays average posterior beliefs about callback rates the HR firm would implement. *Own favela* and *Non-favela* denote a respondent's predicted callback rate for their own favela and adjacent nonfavela, respectively. *Disc* is the implied percent drop in callback rate due to the favela address. The *Actions* panel displays shares of all invited jobseekers in each arm. *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. Sample size in each arm is shown at the bottom of each bar. Vertical error bars display 95% confidence intervals based on robust standard errors, and brackets show p-values for pairwise tests of equality between arms.

Figure 5: Average Effects of Expected Address Visibility on Interview Performance



Note: The graph shows treatment effect estimates (without controls) for the full interview sample (N=422). 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.

A Supporting Tables and Figures

Table A.1: Census (2010) Summary Statistics

Location	Population	Literate Share	White Population Share	Income per Capita in R\$ (2010)	Share Households w/ Per Capita Income $\leq$ MW
All non-favela neighborhoods in Rio	4,888,663	0.92	0.57	1376.35	0.31
All favela neighborhoods in Rio	1,391,953	0.84	0.33	382.87	0.67
Jacarezinho (favela)	37,792	0.87	0.33	349.63	0.71
Manguinhos (favela)	36,151	0.83	0.34	346.86	0.73
Maré (favela)	129,715	0.83	0.38	395.38	0.66
Bonsucesso (non-favela)	18,341	0.93	0.60	897.97	0.31
Maria da Graça (non-favela)	7,967	0.93	0.67	1126.26	0.23

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. MW stands for minimum wage. The first two rows aggregate all favela and all non-favela neighborhoods in the city of Rio de Janeiro; the remaining rows report the study neighborhoods and their adjacent comparison neighborhoods. These are the only Census variables, among those relevant to our study, that are publicly available at the neighborhood level (e.g., average years of schooling is not available). The last column is the share of all households whose monthly per capita income is strictly positive but at most one minimum wage. Households reporting no income are a separate census class; they are included in the denominator but not the numerator.

Table A.2: Expected Discrimination Predictors

	(1) Expects > 50% disc in audit	(2) Expects > 50% disc in audit
Age	0.169 (0.187)	0.059 (0.219)
Male (0/1)	-6.606*** (2.370)	-3.809 (2.841)
White jobseeker (0/1)	-3.634 (2.583)	-4.905 (3.025)
Some college (0/1)	9.333** (3.870)	5.123 (4.677)
Completed regular high-school (0/1)	2.603 (2.827)	0.073 (3.249)
Working now (0/1)	1.418 (2.946)	-4.111 (3.410)
Holds a formal job (0/1)	-5.480 (3.901)	-1.836 (4.602)
Ever worked (0/1)	-1.286 (2.820)	2.276 (3.349)
Actively searched last week (0/1)	2.988 (2.157)	3.040 (2.596)
Reservation wage (USD)	0.024*** (0.009)	0.017* (0.009)
Distance to work (is reason, 0/1)		-0.760 (2.566)
Missing days because of police raids (is reason, 0/1)		5.308* (3.097)
Lower skill (is reason, 0/1)		-0.304 (2.783)
Difficulty adapting to work (is reason, 0/1)		-2.100 (2.818)
Fear or violence (is reason, 0/1)		3.500 (2.726)
Racism (is reason, 0/1)		10.346*** (3.108)
Having a different culture/speech (is reason, 0/1)		4.895* (2.915)
Dislike of favela residents (is reason, 0/1)		3.249 (2.900)
Nepotism (is reason, 0/1)		-5.571** (2.740)
Observations	2166	1496

Note: Multivariate OLS estimates. Outcome is a dummy variable for whether the jobseeker expected at-or-above-median discrimination when predicting the audit study, rescaled so that coefficients are in percentage points. See notes to Table A.3 for independent variable descriptions. Robust standard errors shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: Baseline Statistics

	Mean	SD	Min	Max	N
Maré resident (0/1)	0.62	0.48	0	1	2,167
Jacarezinho resident (0/1)	0.19	0.39	0	1	2,167
Manguinhos resident (0/1)	0.19	0.39	0	1	2,167
Age	25.90	6.25	18	41	2,167
Male (0/1)	0.30	0.46	0	1	2,167
White jobseeker (0/1)	0.22	0.42	0	1	2,167
Some college (0/1)	0.08	0.27	0	1	2,167
Completed regular high-school (0/1)	0.80	0.40	0	1	2,167
Working now (0/1)	0.32	0.47	0	1	2,167
Holds a formal job (0/1)	0.13	0.34	0	1	2,167
Ever worked (0/1)	0.75	0.43	0	1	2,167
Actively searched in the last 7 days	0.49	0.50	0	1	2,167
Surveyor-assessed comm skills (Likert scale, 0-4)	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 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 nine variables are dummies for whether the jobseekers agreed a specific reason was important for explaining why employers might avoid hiring from favelas. Those nine questions were removed after we introduced the information experiment, to control survey duration.

Table A.4: 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)
<i>Favela Info</i>	0.02 (0.05)	0.06 (0.10)	-0.02 (0.13)
<i>Full Info</i>	0.02 (0.05)	-0.12 (0.10)	0.02 (0.14)
Observations	690	389	389
<i>No Info</i> Mean	0.6	2.3	2.5
<i>Favela=Full p</i>	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. None of the specifications include controls. The bottom rows report the mean of the outcome in the *No Info* arm and the p-value for the test that the *Favela Info* and *Full Info* effects are equal. Robust standard errors are shown in parentheses.

Table A.5: Comparison of Samples Across the Three Experiments

Variable	(1) Address Omission Experiment		(2) Information Experiment		(3) Interview Experiment		(1)-(2) Pairwise t-test		(1)-(3) Pairwise t-test		(2)-(3) Pairwise t-test	
	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 (0.011)	690	0.193 (0.015)	422	0.204 (0.020)	1993	-0.009	1725	-0.020	1112	-0.011
Manguinhos resident (0/1)	1303	0.027 (0.004)	690	0.454 (0.019)	422	0.182 (0.019)	1993	-0.427***	1725	-0.156***	1112	0.271***
Age	1303	25.772 (0.174)	690	26.029 (0.236)	422	24.806 (0.283)	1993	-0.257	1725	0.966***	1112	1.223***
Male (0/1)	1303	0.295 (0.013)	690	0.303 (0.018)	422	0.265 (0.022)	1993	-0.008	1725	0.029	1112	0.037
White jobseeker (0/1)	1303	0.229 (0.012)	690	0.210 (0.016)	422	0.237 (0.021)	1993	0.019	1725	-0.008	1112	-0.027
Working now (0/1)	1303	0.326 (0.013)	690	0.284 (0.017)	422	0.135 (0.017)	1993	0.042*	1725	0.191***	1112	0.149***
Holds a formal job (0/1)	1303	0.118 (0.009)	690	0.135 (0.013)	422	0.047 (0.010)	1993	-0.017	1725	0.071***	1112	0.087***
Ever worked (0/1)	1303	0.722 (0.012)	690	0.786 (0.016)	422	0.737 (0.021)	1993	-0.063***	1725	-0.015	1112	0.049*
Actively searched in the last 7 days	1303	0.531 (0.014)	690	0.425 (0.019)	422	0.649 (0.023)	1993	0.106***	1725	-0.118***	1112	-0.225***
High skill (0/1)	1303	0.487 (0.014)	690	0.558 (0.019)	422	0.559 (0.024)	1993	-0.071***	1725	-0.073***	1112	-0.001
Expected discrimination predicting audit results	1303	54.059 (0.823)	690	37.111 (1.281)	422	47.984 (1.673)	1993	16.948***	1725	6.074***	1112	-10.874***
Reservation wage (USD)	1302	253.099 (3.014)	690	246.173 (3.215)	422	231.962 (2.736)	1992	6.927	1724	21.137***	1112	14.211***
F-test of joint significance (F-stat)							86.622***		22.299***		18.542***	
F-test, number of observations							1992		1724		1112	

Notes: Pair-wise comparisons of average baseline characteristics across experiments. The reservation wage statistic drops the one respondent with a missing reservation wage. See notes to Table A.3 for details on variables. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.6: Pearson Correlation Coefficients Between Interview Performance Indicators

	Overall (Interv.)	Calm (Interv.)	Professional (Interv.)	Overall (Self)	Calm (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		
Calm (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 performance indexes, and "Interv." refers to interviewer assessments. Standard errors in parentheses.

Table A.7: Address Omission Experiment: Randomization Balance

Variable	(1) Address Omission		(2) Known Address		(3) Status Quo		(1)-(2)		(1)-(3) Pairwise t-test		(2)-(3)	
	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	25.894 (0.301)	437	26.085 (0.303)	432	25.333 (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	54.034 (1.425)	437	53.867 (1.441)	432	54.277 (1.412)	871	0.167	866	-0.243	869	-0.410
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)							1.565*		1.372		0.776	
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. Expected discrimination when predicting the audit results is the implied percent drop in callback rates if a résumé lists a favela address (winsorized at the first and 99th percentiles). See notes to Table A.3 for other variable descriptions. * p<0.1, ** p<0.05, *** p<0.01.

Table A.8: Information Experiment: Randomization Balance

Variable	(1) Favela Info		(2) Full Info		(3) No Info		(1)-(2)		(1)-(3) Pairwise t-test		(2)-(3)	
	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	26.773 (0.412)	227	25.463 (0.394)	216	25.773 (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	38.052 (2.117)	227	34.835 (2.326)	216	38.426 (2.220)	474	3.216	463	-0.374	443	-3.591
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)							0.911		0.991		0.704	
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. Expected discrimination when predicting the audit results is the implied percent drop in callback rates if a résumé lists a favela address (winsorized at the first and 99th percentiles). See notes to Table A.3 for other variable descriptions. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.9: Interview Experiment: Randomization Balance

Variable	(1) Name-Only Mean/(SE)	(2) Name-and-Address Mean/(SE)	(1)-(2) Pairwise t-test Mean difference
<i>Address Omission</i>	0.227 (0.028)	0.243 (0.030)	-0.015
<i>Known Address</i>	0.218 (0.028)	0.233 (0.030)	-0.014
<i>Favela Info</i>	0.141 (0.024)	0.153 (0.025)	-0.013
<i>Full Info</i>	0.150 (0.024)	0.134 (0.024)	0.016
Jacarezinho resident (0/1)	0.232 (0.029)	0.173 (0.027)	0.059
Manguinhos resident (0/1)	0.200 (0.027)	0.163 (0.026)	0.037
Age	24.768 (0.409)	24.847 (0.390)	-0.078
Male (0/1)	0.268 (0.030)	0.262 (0.031)	0.006
White jobseeker (0/1)	0.236 (0.029)	0.238 (0.030)	-0.001
Some college (0/1)	0.064 (0.016)	0.079 (0.019)	-0.016
Completed regular high-school (0/1)	0.786 (0.028)	0.767 (0.030)	0.019
Working now (0/1)	0.164 (0.025)	0.104 (0.022)	0.060*
Holds a formal job (0/1)	0.050 (0.015)	0.045 (0.015)	0.005
Ever worked (0/1)	0.764 (0.029)	0.708 (0.032)	0.056
Actively searched last week (0/1)	0.627 (0.033)	0.673 (0.033)	-0.046
Skill index	0.032 (0.037)	0.009 (0.040)	0.023
Expected discrimination predicting audit results	48.372 (2.294)	47.562 (2.451)	0.810
Reservation wage (USD)	232.233 (3.830)	231.667 (3.918)	0.565
F-test of joint significance (F-stat)			0.819
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. Expected discrimination when predicting the audit results is the implied percent drop in callback rates if a résumé lists a favela address (winsorized at the first and 99th percentiles). See notes to Table A.3 for other variable descriptions. * p<0.1, ** p<0.05, *** p<0.01

Table A.10: Across Experiments, Only the Heterogeneity by Race is Robust

	Experiment-specific outcomes			
	(1)	(2)	(3)	(4)
Agg. interv. performance (SD) <i>Name-Only</i> $\times h_i$	0.24* (0.14)	0.24** (0.12)	0.07 (0.13)	-0.12 (0.11)
Show-up (%) <i>Status Quo</i> $\times h_i$	-7.22 (6.70)	3.89 (5.55)	2.18 (5.59)	1.17 (5.42)
Show-up (%) <i>Known Address</i> $\times h_i$	-15.52** (6.37)	-1.77 (5.52)	7.33 (5.62)	-5.69 (5.26)
Show-up (%) <i>Favela Info</i> $\times h_i$	-0.23 (9.81)	2.61 (7.66)	0.29 (7.99)	14.37* (7.57)
Show-up (%) <i>Full Info</i> $\times h_i$	-18.00** (8.93)	-5.99 (7.61)	-2.89 (7.89)	11.80 (7.59)
Heterogeneity variable h	White	High $\mathbb{E}[\text{disc}]$	Male	High Skill
Any heterogeneity by h , p-value	0.01	0.241	0.826	0.191
Heterogeneity by h in $\mathbb{E}[\text{address visib.}]$ treatments, p-value	0.023	0.142	0.574	0.363
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 at once. Each column presents the regression coefficients from a stacked regression in which each layer of the stack is akin to 3, so each displayed coefficient is numerically the same as the difference between the corresponding within-group effects in the per-experiment tables. Standard errors and p-values are based on a clustered variance-covariance matrix. The reported p-values can be Bonferroni-corrected by multiplying them by four.

Table A.11: Heterogeneity in Address Omission Experiment Including Interacted Covariates

	(1) Show Up (%)	(2) Show Up (%)	(3) Show Up (%)	(4) Show Up (%)
<i>Status Quo</i> $\times h$	-4.72 (6.14)	1.32 (3.40)	0.53 (4.60)	0.55 (4.37)
<i>Known Address</i> $\times h$	-12.83** (5.81)	-1.74 (3.33)	1.94 (4.84)	-1.21 (4.22)
<i>Status Quo</i> $\times (1 - h)$	0.83 (3.05)	-3.78 (4.45)	-0.84 (3.31)	-1.08 (4.05)
<i>Known Address</i> $\times (1 - h)$	1.00 (3.01)	-2.89 (4.48)	-3.88 (3.21)	-3.97 (3.99)
<i>h</i>	8.39* (5.05)	-0.42 (4.00)	-5.54 (3.87)	3.50 (4.75)
Observations	1302	1302	1302	1302
<i>h</i>	White	High E[disc.]	Male	More skilled
<i>Address Omission</i> mean when $h = 0$	17.18	18.92	22.42	18.50
<i>Status Quo</i> = <i>Known Address</i> when $h = 1$, p-val	0.15	0.36	0.78	0.69
<i>Status Quo</i> = <i>Known Address</i> when $h = 0$, p-val	0.95	0.84	0.33	0.42
Effect of <i>Status Quo</i> constant across h , p-val	0.43	0.36	0.81	0.80
Effect of <i>Known Address</i> constant across h , p-val	0.04	0.84	0.32	0.66

Note: OLS estimates for the effects of *Status Quo* and *Known Address* on show-up rates for groups with $h_i = 1$ and $h_i = 0$. Regressions include the baseline covariates from Table A.7, both in levels and interacted with treatments. The scalars report p-values for tests of equality: whether treatment effects have equal effects within each group (rows 3-4), and whether treatment effects differ by h (rows 5-6). Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.12: Heterogeneity in Information Experiment Including Interacted Covariates

	(1) Show Up (%)	(2) Show Up (%)	(3) Show Up (%)	(4) Show Up (%)
<i>Full Info</i> $\times h$	-17.17** (8.30)	-4.55 (6.02)	-6.89 (6.21)	4.15 (5.62)
<i>Favela Info</i> $\times h$	-3.34 (8.99)	-1.75 (6.03)	-4.49 (6.63)	3.69 (5.85)
<i>Full Info</i> $\times (1 - h)$	2.85 (4.26)	1.11 (4.79)	1.05 (4.66)	-9.12 (5.80)
<i>Favela Info</i> $\times (1 - h)$	0.94 (4.28)	1.31 (5.09)	2.01 (4.83)	-3.80 (5.82)
<i>h</i>	6.78 (7.42)	5.45 (5.86)	-2.20 (6.07)	5.29 (6.66)
Observations	690	690	690	690
<i>h</i>	White	High E[disc.]	Male	More skilled
<i>No Info</i> mean when $h = 0$	18.60	18.33	21.43	21.59
<i>Full = Favela</i> when $h = 1$, p-val	0.07	0.61	0.66	0.93
<i>Full = Favela</i> when $h = 0$, p-val	0.64	0.97	0.84	0.29
Effect of <i>Full Info</i> constant across h , p-val	0.03	0.46	0.31	0.13
Effect of <i>Favela Info</i> constant across h , p-val	0.67	0.70	0.44	0.40

Note: OLS estimates for the effects of *Full Info* and *Favela Info* on show-up rates for groups with $h_i = 1$ and $h_i = 0$. Regressions include the baseline covariates from Table A.8, both in levels and interacted with treatments. The scalars report p-values for tests of equality: whether treatment effects have equal effects within each group (rows 3-4), and whether treatment effects differ by h (rows 5-6). Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.13: Heterogeneity in the Interview Experiment Including Interacted Covariates

	(1) Aggregate performance index	(2) Interviewer- assessed performance	(3) Self- assessed performance	(4) Aggregate performance index	(5) Interviewer- assessed performance	(6) Self- assessed performance	(7) Aggregate performance index	(8) Interviewer- assessed performance	(9) Self- assessed performance	(10) Aggregate performance index	(11) Interviewer- assessed performance	(12) Self- assessed performance
<i>Name-Only</i> $\times h$	0.31*** (0.12)	0.32** (0.16)	0.30** (0.13)	0.18** (0.07)	0.17* (0.10)	0.19** (0.08)	0.15 (0.11)	0.04 (0.16)	0.27** (0.13)	0.11 (0.09)	0.12 (0.12)	0.10 (0.10)
<i>Name-Only</i> $\times (1 - h)$	0.03 (0.06)	-0.03 (0.09)	0.09 (0.08)	-0.03 (0.09)	-0.12 (0.12)	0.05 (0.10)	0.08 (0.07)	0.07 (0.09)	0.09 (0.08)	0.08 (0.11)	-0.02 (0.15)	0.18 (0.12)
<i>h</i>	-0.08 (0.10)	0.01 (0.13)	-0.16 (0.11)	-0.04 (0.08)	-0.09 (0.11)	0.02 (0.10)	0.01 (0.10)	0.00 (0.13)	0.02 (0.12)	0.06 (0.11)	0.12 (0.16)	0.01 (0.13)
Observations	422	422	422	422	422	422	422	422	422	422	422	422
<i>h</i>	White	White	White	High E[disc.]	High E[disc.]	High E[disc.]	Male	Male	Male	More skilled	More skilled	More skilled
<i>Name-and-Address</i> mean when $h = 0$	.04	.01	.06	.04	.07	.01	.02	.02	.02	-.02	-.05	.02
Effect is constant across h , p-val	0.04	0.06	0.18	0.06	0.06	0.30	0.56	0.87	0.23	0.88	0.55	0.63

Note: OLS estimates for the effects of *Name-Only* on interview performance indexes for groups with $h_i = 1$ and $h_i = 0$. Sample restricted to those who showed up to the interview. Regressions include the baseline covariates from Table A.9, both in levels and interacted with the treatment. The scalar reports the p-value for testing whether the treatment effect differs by h . Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.14: Obfuscation Predictors

	(1) Obfuscates in app (%)	(2) Obfuscates at office (%)	(3) Obfuscates in app (%)	(4) Obfuscates at office (%)
Age	0.93** (0.41)	0.09 (0.46)	-0.25 (0.62)	-0.36 (0.58)
Male (0/1)	1.05 (4.83)	8.27 (5.07)	6.71 (8.11)	6.40 (7.22)
White jobseeker (0/1)	-2.48 (4.84)	-1.48 (5.06)	-5.58 (7.22)	-6.58 (6.71)
Some college (0/1)	0.62 (8.00)	-7.69 (7.59)	-4.67 (12.71)	-3.00 (11.58)
Completed regular high-school (0/1)	-4.70 (5.44)	-2.83 (5.45)	-0.10 (7.82)	-1.99 (7.53)
Working now (0/1)	17.80** (7.01)	2.80 (8.10)	7.40 (13.30)	17.88 (13.16)
Holds a formal job (0/1)	-10.90 (10.20)	-12.22 (11.67)	2.53 (22.27)	-33.99** (17.11)
Ever worked (0/1)	-4.75 (5.02)	-3.28 (5.53)	2.35 (7.86)	1.33 (7.54)
Actively searched last week (0/1)	3.54 (4.16)	0.18 (4.43)	2.81 (6.69)	0.60 (6.20)
Reservation wage (USD)	-0.02 (0.03)	0.04 (0.04)	0.08 (0.06)	0.09* (0.05)
Expected discrimination	0.17*** (0.06)	0.14** (0.06)	0.19** (0.08)	0.23*** (0.07)
Observations	502	422	227	227
Dep. var. avg.	27.69	23.22	26.87	21.59
<i>Status Quo</i> app	Yes	No	Yes	Yes
Cond. on being interviewed	No	Yes	Yes	Yes

Note: Multivariate OLS estimates. Outcome is a dummy variable for whether the jobseeker obfuscated their address in the *Status Quo* application form (Columns 1 and 3) and at the interview office (Columns 2 and 4). Expected discrimination is the posterior expected discrimination when available and baseline expected discrimination when the posterior is not available (win-sorized at the first and 99th percentiles); see notes to Table A.3 for other independent variable descriptions. Robust standard errors shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.15: No Effects on Show-Up By Names Distinctive of Low-SES

	(1) Applied (%)	(2) Applied (%)	(3) Show Up (%)	(4) Show Up (%)
<i>AddressOmission</i> × nonlow SES name	-6.25 (4.13)	-7.36* (4.09)	1.74 (3.20)	1.01 (3.21)
<i>AddressOmission</i> × low SES name	-1.13 (5.82)	1.53 (5.47)	4.47 (4.66)	5.26 (4.47)
<i>StatusQuo</i> × nonlow SES name	-1.71 (4.19)	-2.65 (4.05)	4.69 (3.32)	4.09 (3.28)
<i>StatusQuo</i> × low SES name	-4.57 (5.74)	-2.18 (5.51)	-2.88 (4.26)	-1.34 (4.16)
Observations	1303	1302	1303	1302
Controls	No	Yes	No	Yes

Notes: Table displays OLS estimates of treatment effects on the application outcome (columns 1 and 2) and interview attendance (columns 3 and 4). Controls by type of name are always included, and all variables in Table A.7 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. Specifications including controls (even columns) drop the one respondent with a missing reservation wage. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.16: 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 (%)	-0.53 (0.56)	-0.35 (0.47)	-0.70* (0.37)
Posterior Expected Discrimination Rate (%)	-0.06 (0.24)	0.20 (0.19)	0.16 (0.17)
Observations	690	690	690
<i>No Info</i> Mean	39.8	19.9	10.2

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.17: No Effects on Interview Performance By Names Distinctive 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
<i>Name-Only</i> × low SES name	0.05 (0.11)	0.04 (0.11)	-0.02 (0.15)	-0.04 (0.15)	0.12 (0.12)	0.12 (0.12)
<i>Name-Only</i> × nonlow SES name	0.16** (0.07)	0.11* (0.07)	0.13 (0.10)	0.08 (0.09)	0.19** (0.08)	0.14* (0.08)
Observations	422	422	422	422	422	422
p-value: equal effects across name groups	0.38	0.57	0.38	0.50	0.62	0.87
Controls	No	Yes	No	Yes	No	Yes

Notes: Table displays OLS estimates of the effect of *Name-Only* on interview performance within each first-name group; each coefficient is the effect relative to *Name-and-Address* for the same group, with no omitted heterogeneity category. The “p-value: equal effects across name groups” row tests equality of the two coefficients. Even columns include the baseline covariates from Table A.9, demeaned and interacted with the treatment. 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.18: Treatment Effects on Performance by Obfuscation at Application and Interview Stage

	(1) Aggregate performance index	(2) Interviewer-assessed performance index	(3) Self-assessed performance index
<i>Name-Only</i> × <i>S2obf</i> × <i>S3rev</i> (<i>N</i> =18)	0.28 (0.28)	0.20 (0.41)	0.35 (0.23)
<i>Name-Only</i> × <i>S2rev</i> × <i>S3obf</i> (<i>N</i> =6)	-1.40 (1.42)	-1.58 (1.56)	-1.22 (1.13)
<i>Name-Only</i> × <i>S2rev</i> × <i>S3rev</i> (<i>N</i> =161)	0.20** (0.09)	0.22* (0.12)	0.19* (0.11)
<i>Name-Only</i> × <i>S2obf</i> × <i>S3obf</i> (<i>N</i> =43)	0.26 (0.19)	0.43* (0.24)	0.08 (0.22)
<i>Name-Only</i> × <i>KA</i> × <i>S3 obf</i> (<i>N</i> =29)	0.01 (0.26)	0.12 (0.33)	-0.11 (0.31)
<i>Name-Only</i> × <i>KA</i> × <i>S3 rev</i> (<i>N</i> =66)	0.13 (0.13)	-0.08 (0.18)	0.34** (0.17)
<i>Name-Only</i> × <i>AO</i> × <i>S3 obf</i> (<i>N</i> =20)	0.12 (0.36)	0.08 (0.56)	0.17 (0.34)
<i>Name-Only</i> × <i>AO</i> × <i>S3 rev</i> (<i>N</i> =79)	-0.05 (0.14)	-0.26 (0.21)	0.15 (0.13)
Observations	422	422	422
Effect on those who obfuscate twice equals the effect on those who reveal twice, <i>p</i> -value	0.79	0.43	0.67


Note: OLS estimates of the effects of *Name-Only* on interview performance indexes by joint obfuscation behavior at the application (Stage 2) and interview (Stage 3) stages. *S2* and *S3* refer to obfuscation decisions at the application and interview stages, respectively; *obf* indicates obfuscation and *rev* indicates true revelation of address. *KA* (*Known Address*) and *AO* (*Address Omission*) denote application-stage treatments in which obfuscation was mechanically constrained. The sample is restricted to interview attendees. Coefficients are in standard deviation units of interview performance. The final row reports *p*-values for tests of equality between candidates who always obfuscate and those who always reveal their address. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.19: Effects on Interview Performance 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
<i>Name-Only</i> × <i>Status Quo</i> × <i>No Info</i> (N=129)	0.23** (0.10)	0.12 (0.10)	0.35** (0.14)	0.25* (0.13)	0.10 (0.12)	-0.01 (0.12)
<i>Name-Only</i> × <i>Status Quo</i> × <i>Favela Info</i> (N=55)	0.25 (0.18)	0.30* (0.17)	0.33 (0.23)	0.37* (0.21)	0.16 (0.20)	0.22 (0.20)
<i>Name-Only</i> × <i>Status Quo</i> × <i>Full Info</i> (N=44)	0.06 (0.16)	0.07 (0.16)	-0.18 (0.23)	-0.13 (0.23)	0.30* (0.18)	0.27 (0.18)
<i>Name-Only</i> × <i>Address Omission</i> × <i>No Info</i> (N=90)	-0.08 (0.13)	-0.07 (0.12)	-0.28 (0.20)	-0.27 (0.19)	0.12 (0.13)	0.12 (0.12)
<i>Name-Only</i> × <i>Known Address</i> × <i>No Info</i> (N=81)	0.09 (0.13)	0.00 (0.14)	0.07 (0.18)	-0.05 (0.18)	0.10 (0.16)	0.06 (0.16)
<i>Name-Only</i> × <i>Other Combinations</i> (N=23)	0.25 (0.25)	0.34 (0.25)	-0.11 (0.41)	-0.03 (0.41)	0.61** (0.27)	0.71** (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	0.50	0.53
P-value for no effect on less externally valid subsample	0.75	0.66	0.59	0.65	0.05	0.06
P-value for equal effect on more vs. less externally valid	0.10	0.19	0.00	0.01	0.55	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 in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure A.1: 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. Estimates use no controls, as in the other heterogeneity break-downs. See notes to Figure 4 for details on outcomes and figure features.

Figure A.2: 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.3: 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]”.

Figure A.4: Conditional Treatment Effects – Address Omission Experiment




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 3 for details and Section 4.1 for a systematic discussion of treatment effect heterogeneity.

Figure A.5: 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 4 for details and Section 4.2 for a systematic discussion of treatment effect heterogeneity.

Figure A.6: 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 4 for details and Section 4.2 for a systematic discussion of treatment effect heterogeneity.

Figure A.7: Conditional Treatment Effects of *Name-Only*




Notes: Preregistered heterogeneity break-downs for the interview experiment. See notes to Figure 5 for details and Section 4.3 for a systematic discussion of treatment effect heterogeneity.

Figure A.8: Effects of Information Treatments on Beliefs and Applications by Whether Job-seekers Initially Under- or Overestimated the Favela Callback Rate




Note: Same as in Figure 4, 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.9: 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.

Figure A.10: Figure 3 with Lasso-selected Controls




Notes: The double-lasso procedure could pick controls including expected discrimination, race, gender, residence in Jacarezinho and Manguinhos, age, having ever worked, working now, whether the jobseeker has looked for a job in the last 7 days, an indicator for being more skilled, and the reservation wage. After picking controls, we demean them and interact them with treatment status, so the bar heights reflect control-corrected means and differences reflect ATE between arms. The variance-covariance matrix was estimated using the HC1 approach since the lasso-selected controls can generate high-leverage observations; see notes for Figure 3 for other details.

Figure A.11: Figure 4 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. The variance-covariance matrix was estimated using the HC1 approach since the lasso-selected controls can generate high-leverage observations; see notes for Figure 4 for other details.

Figure A.12: Figure 5 with Lasso-Selected Controls



Note: The graph shows treatment effect estimates (using double-lasso selected controls) for the full sample. The double-lasso procedure could pick controls including expected discrimination, race, gender, residing in Manginhos or Jacarezinho, 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.

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 heterogeneity. Below, we list our deviations from that plan.

- 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 a primary outcome. For brevity, we omit them from the main text and present them in this section.
- 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 new interviewees was too low for making 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 candidate flow into interviews, the batch sizes had to be smaller to keep such flow constant.

- 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).
- We piloted the Information Experiment with partial data from the audit study, which for some callbacks, some which came in late and some that could not be matched to a single specific job (but could still be matched to a name based on contact information). During the pilot, the Maré callback rate was 15%, and the Bonsucesso callback rate was 16.5%. We conducted the pilot among N=150 participants while still running the Address Omission Experiment. We explain how we use these observations in the next bullet.
- We randomly assigned 183 out of the main sample of 2,167 jobseekers to a condition in the Address Omission Experiment and to a condition in the Information Experiment. Out of those 183, 150 participated in the aforementioned pilot of the Information Experiment. The remaining 33 participated in the final version of the Information Experiment, as described in the main text. We included the 10 participants who were randomly assigned to *Status Quo* in the main sample for the Information Experiment, since their application procedure matches the one used in the main sample. There was also **one** additional person who did not participate in the Address Omission Experiment but participated in the Information Experiment pilot, whom we nevertheless invited to apply to the jobs but who we do not include in the main sample of the Information Experiment. To summarize, **N=2,167** sample consists of: **N=1,303** who received only Address Omission assignments; **N=690** who received only Information assignments or received the *Status Quo* assignment in the overlapping sample of the Information Experiment and Address Omission Experiment; **N=174** whom we do not include in the main sample for the Information or Address Omission Experiment but we invited to apply to our partner's jobs, so they may be included in the Interview Experiment. The overlapping sample analysis thus adds 183 observations to the Address Omission Experiment and 184 to the Information Experiment samples. Appendix B.2 reproduces our main tables and figures including the overlapping

sample.

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)
<i>Address Omission</i>	-1.44 (3.32)	5.13 (7.98)	0.65 (1.08)	0.15 (0.10)	-0.00 (0.01)	0.03 (0.04)	-0.07 (0.06)	-0.05 (0.07)	-0.07 (0.07)
<i>Known Address</i>	-0.25 (3.31)	7.09 (7.35)	0.45 (1.03)	0.04 (0.09)	-0.00 (0.01)	0.00 (0.04)	-0.08 (0.06)	-0.04 (0.07)	-0.03 (0.07)
Observations	1303	1303	1303	1303	1303	975	422	341	341
<i>Status Quo Average</i>	61.57	52.43	8.32	0.46	0.03	0.42	0.40	0.64	0.63
<i>Address Omission=KnownAddress p</i>	0.72	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 Information Experiment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
	Words in cover letter	Cover letter quality (0/100)	Exp. Disc. on the job (% chance)	Exp. disc for self (%)	Wage gap for self (%)	# of planned future applications	Exp. future callback rate (%)	More excited about job search (0/1)	Agrees job mkt is unfair (-2/2)	Agrees job search is an opportunity (-2/2)	Agrees one must ignore disc. (-2/2)	Agrees one must apply to all jobs (-2/2)	Plans to increase search effort (0/1)	Feels is in bad position by endline (1/4)	Improved search strategy by endline (0/1)
<i>Favela Info</i>	15.80 (10.34)	1.50 (1.57)	-6.47** (2.54)	-2.77 (3.09)	-1.26 (2.95)	-8.27 (6.46)	-2.74 (2.66)	-0.22*** (0.04)	-0.01 (0.12)	-0.13 (0.11)	0.21* (0.11)	-0.04 (0.13)	-0.02 (0.03)	-0.04 (0.11)	-0.10 (0.06)
<i>Full Info</i>	15.41 (10.36)	0.33 (1.54)	-5.29** (2.66)	-14.63** (6.62)	-9.29 (8.59)	-3.37 (6.92)	0.72 (2.78)	-0.15*** (0.04)	-0.13 (0.13)	0.09 (0.11)	0.00 (0.11)	0.10 (0.13)	0.01 (0.03)	0.14 (0.11)	-0.04 (0.06)
Observations	690	690	690	689	690	690	690	670	670	670	670	670	690	389	389
<i>No Info</i> 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 <i>p</i>	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 against the same counterfactual 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 cannot 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.3: 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)
<i>Name-Only</i>	0.00 (0.03)	0.03 (0.02)	-0.01 (0.01)	0.08 (0.08)	-0.07 (0.11)	0.07 (0.09)	0.07 (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 interview; (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 (7) had little variation, about 85% of candidates picked ten out of ten in those questions.

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

In this Section, we present the main tables and figures adding 183 participants between the two experiments who participated in the Address Omission and Information Experiments.

Table B.4: Table A.4 Including Overlapping Sample

	(1) Responded to endline (0/1)	(2) Exp. discrimination (categorical, 1-4)	(3) # of sent apps (categorical, 1-4)
<i>Favela Info</i>	-0.00 (0.04)	0.02 (0.08)	-0.03 (0.12)
<i>Full Info</i>	0.01 (0.04)	-0.15* (0.09)	0.06 (0.12)
Observations	863	505	507
<i>No Info</i> Mean	0.6	2.3	2.5
Favela=Full <i>p</i>	0.75	0.04	0.47

Note: See notes for Table A.4 for details.

Table B.5: Heterogeneity in Address Omission Experiment Including Overlapping Sample

	(1) Show Up (%)	(2) Show Up (%)	(3) Show Up (%)	(4) Show Up (%)
<i>Status Quo</i> $\times h$	-5.68 (5.63)	-0.26 (3.23)	0.30 (4.15)	-1.43 (3.71)
<i>Known Address</i> $\times h$	-12.86** (5.36)	-2.10 (3.18)	7.42 (4.52)	-4.12 (3.58)
<i>Status Quo</i> $\times (1 - h)$	-0.09 (2.81)	-3.39 (4.00)	-2.44 (3.15)	-0.71 (3.43)
<i>Known Address</i> $\times (1 - h)$	2.59 (2.86)	1.34 (4.21)	-4.51 (3.08)	2.48 (3.58)
<i>h</i>	8.83* (4.59)	2.40 (3.72)	-7.31** (3.60)	4.83 (4.38)
Observations	1486	1486	1486	1486
<i>h</i>	White	High $\mathbb{E}[disc.]$	Male	More skilled
<i>Address Omission</i> mean when $h = 0$	17.89	18.29	22.46	17.33
<i>Status Quo</i> = <i>Known Address</i> when $h = 1$, p-val	0.16	0.56	0.13	0.47
<i>Status Quo</i> = <i>Known Address</i> when $h = 0$, p-val	0.35	0.24	0.48	0.35
Effect of <i>Status Quo</i> constant across h , p-val	0.37	0.54	0.60	0.89
Effect of <i>Known Address</i> constant across h , p-val	0.01	0.51	0.03	0.19


Note: OLS estimates for the effects of *Status Quo* and *Known Address* on show-up rates for groups with $h_i = 1$ and $h_i = 0$. Regressions control for male and skill index. The scalars report p-values for tests of equality: whether treatment effects have equal effects within each group (rows 3-4), and whether treatment effects differ by h (rows 5-6). Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.6: Heterogeneity in Information Experiment Including Overlapping Sample

	(1) Show Up (%)	(2) Show Up (%)	(3) Show Up (%)	(4) Show Up (%)
<i>Full Info</i> $\times h$	-9.09 (6.90)	0.27 (5.19)	-4.09 (6.16)	5.81 (4.77)
<i>Favela Info</i> $\times h$	-0.49 (7.36)	1.76 (5.02)	-3.18 (5.96)	5.07 (4.59)
<i>Full Info</i> $\times (1 - h)$	4.62 (3.94)	2.97 (4.61)	4.20 (4.15)	-3.51 (4.93)
<i>Favela Info</i> $\times (1 - h)$	1.21 (3.73)	-0.02 (4.45)	2.68 (3.99)	-4.59 (4.76)
<i>h</i>	2.80 (5.79)	2.74 (4.76)	2.23 (5.26)	1.04 (4.82)
Observations	864	864	864	864
<i>h</i>	White	High E[disc.]	Male	More skilled
<i>No Info</i> mean when $h = 0$	18.89	18.24	19.19	19.67
<i>Full</i> = <i>Favela</i> when $h = 1$, p-val	0.22	0.77	0.87	0.88
<i>Full</i> = <i>Favela</i> when $h = 0$, p-val	0.38	0.51	0.72	0.81
Effect of <i>Full Info</i> constant across h , p-val	0.08	0.70	0.26	0.17
Effect of <i>Favela Info</i> constant across h , p-val	0.84	0.79	0.41	0.15


Note: OLS estimates for the effects of *Full Info* and *Favela Info* on show-up rates for groups with $h_i = 1$ and $h_i = 0$. Regressions control for Manguinhos and age. The scalars report p-values for tests of equality: whether treatment effects have equal effects within each group (rows 3-4), and whether treatment effects differ by h (rows 5-6). Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure B.1: Figure 3 Including Overlapping Sample



Note: See notes to Figure 3 for details.

Figure B.2: Figure 4 Including Sample Overlapping with Address Omission Experiment



Note: See notes to Figure 4 for details.

C Audit Study

C.1 Relevance of the Jobs in the Audit to those in the Supply-Side Experiments

Geography. Among jobseekers participating in our baseline survey who have ever worked (N=1,623), 55% reported having worked in the North Zone, 41% in the South Zone, and 39% Downtown, with 80% in at least one of these three areas. In the audit study, Table C.1 shows that 47% of ads declared a location in the North Zone, 21.4% in the South Zone, and 18% Downtown (78% in at least one of these three). Hence, there is a strong overlap between the locations of the audit study jobs and where jobseekers in our sample work.

Occupation. We compare the jobs in the audit study with the last position held by jobseekers in our sample (N=788, since we only ask that question to those whose last job was a formal job). We see meaningful overlap: 34.9% of respondents' last jobs were sales-related (e.g., salesperson, store clerk, telemarketing operator), thus directly comparable to those in the audit study. The remaining most common occupations among jobseekers in our study included general labor (8.5%), administrative assistant (7.7%), receptionist (4.1%), and apprenticeships (2.7%).

C.2 Methodology

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 looked for job posting no older than two weeks on 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 discarded it.

Randomization. We created 8 résumés (4 male, 4 female) with similar levels of experience,

skills and education. Within each gender, two versions were the same, except for the name, contact information, and address (Maré vs. Bonsucesso). Hence, we had 4 unique résumé templates. For each job posting, we randomly ordered the 4 unique templates, then randomly assigned a favela address to one of the first two résumés and a nonfavela address to one of the other two résumés, ensuring the job received two applications from different templates. The remaining two templates served as backups for gender-specific postings. If a job was 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.

Names and socioeconomic status. We present callback rates by each name we used in the audit study in Figure C.1. Each couple of bars belongs to an applicant’s “profile,” meaning that they share all characteristics except for name, address, and contact information. We randomly picked these names among the top-50 most common in Rio’s labor force, so they are widely used and unlikely to convey strong signals about race or SES. The only one among these names with a marker of lower SES is Robson, since it includes the suffix “SON.”³⁹ Applications with the name Robson received a lower callback rate—6.5pp less than Guilherme, its pair, with $p=0.08$ —which could reflect perceptions of lower SES. If we instead drop the Robson/Guilherme pair, we estimate a (nonsignificant) *pro*-favela bias of 1.4pp ($p=0.42$), which does not threaten our initial interpretation of the results.

C.3 Results

³⁹Scottini (2011) uses administrative data on Brazilian students to study how names predict low-SES. The study identifies several features that predict lower SES, such as (i) low frequency in the population, (ii) unusual translation of an English-language name, (iii) featuring the letters K, W, or Y, (iv) ending in “on” or “son”, and others. Notably for our purposes, Scottini (2011) finds a stronger relationship between first names and socioeconomic status than between first names and race.

Table C.1: Summary Statistics on Audit Study Jobs - Location and Industry

	Mean	SD	Min	Max	N
North Zone job	0.47	0.50	0	1	1,400
South Zone job	0.21	0.41	0	1	1,400
Downtown job	0.18	0.39	0	1	1,400
Further South	0.30	0.46	0	1	1,400
More than one location	0.15	0.35	0	1	1,400
Formal contract with employer	0.80	0.40	0	1	1,400
Home or hybrid	0.03	0.17	0	1	1,400
Temporary	0.01	0.11	0	1	1,400
Apparel firm	0.14	0.35	0	1	1,400
Automobiles, machine parts	0.04	0.19	0	1	1,400
Beauty/cosmetics	0.07	0.26	0	1	1,400
Education	0.04	0.19	0	1	1,400
Financial, insurance, real state	0.09	0.28	0	1	1,400
Clinics or drugstores	0.06	0.23	0	1	1,400
Confidential firm	0.05	0.21	0	1	1,400


Note: This table reports summary statistics for the 1,400 audit study applications (700 job postings). Each row indicates whether a job posting belongs to a given location, contract type, or industry category. Jobs may list more than one location. “Confidential firm” indicates postings where the employer name was not disclosed.

Table C.2: Audit Study Results

	(1) Callback (%)	(2) Callback (%)	(3) Callback (%)	(4) Callback (%)	(5) Callback (%)
Maré résumé	-0.34 (1.28)	-0.40 (1.29)	-1.04 (1.18)	0.39 (2.07)	-2.27 (1.66)
Maré résumé × Downtown job				-2.50 (3.98)	
Maré résumé × South Zone job				3.42 (4.19)	
Maré résumé × Further South jobs				-3.02 (3.39)	
Maré résumé × Apparel					5.44 (4.47)
Maré résumé × Beauty/cosmetics					4.98 (5.31)
Maré résumé × Financial services					3.86 (4.03)
Maré résumé × Clinics or drugstores					3.08 (5.30)
Maré résumé × Confidential firm					5.16 (3.75)
Observations	1400	1400	1400	1174	1400
Non-favela Mean	16.96	16.96	16.96	16.75	16.96
Controls	No	Yes	Yes	Yes	Yes
Job FEs	No	No	Yes	No	No


Note: OLS estimates of the effect of having a favela address on the percentage of applications receiving a callback. 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, the hiring regime, and résumé template; columns 2 and 5 additionally include industry-category dummies. In column 4, the sample is restricted to jobs that coders could trace back to a single location out of four possibilities: (i) North Zone (the omitted category), which includes Maré and Bonsucesso, (ii) Downtown, where we had our interview office, (iii) South Zone, which is considerably richer than North Zone, and (iv) Further south, which is also richer but further away. In column 5, the omitted category stands for firms that could not be grouped into categories larger than 5%. Financial services include selling insurance policies, real state brokerage, and loans. The callback level in this table is ≈ 3 pp 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 in parentheses.

Figure C.1: Audit Study Callback Rates by Name



Note: This figure plots callback rates by applicant name and address type in the audit study. For each name, bars compare applications listing a favela address (Maré) and a nonfavela address (Bonsucesso). Each pair of names corresponds to résumés that are otherwise comparable in qualifications, experience, and layout. Vertical lines denote 95% confidence intervals.

Figure C.2: Satellite Image: Bonsucesso (nonfavela) and Maré (Favela)



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

Figure C.3: Example Résumé – Maré home address

ROBSON DE FREITAS

30 YEARS OLD • BRAZILIAN • SINGLE

CONTACT

(21) 99878-2186

defreitasrobson1@gmail.com

Carlos Lacerda Street, 102
- Maré, RJ

EDUCATION

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

SENAC
Logistics Technician.
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.4: Example Résumé – Bonsucesso Address

GUILHERME ANTÔNIO ALMEIDA

30 YEARS OLD • BRAZILIAN • SINGLE

CONTACT

 (21) 99878-2186

 guilhermeantonioalmeida3@gmail.com

 João Torquato Street, 133
- Bonsucesso, RJ

EDUCATION

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

SENAC
Logistics Technician.
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 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 someone 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.3 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 MARÉ's 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 Address Omission Experiment - 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/lyhWIDS5*

D.3 Application Form

Figure D.2: Second Screen of the Application Form—Experimental Condition in the Address Omission Experiment

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

Please confirm your NAME:

☐ My name is {e://Field/name}

☐ No, my correct NAME is:

Please confirm your MOBILE PHONE/WHATSAPP:

☐ My MOBILE PHONE/WHATSAPP is the one in which I received the link to this form

☐ No, my correct MOBILE PHONE/WHATSAPP is:

We also need your home address:

Street

Number and unit if applicable

Neighborhood -- start typing and select your neighborhood

(a) *Status Quo*



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

Please confirm your NAME:

☐ My name is {e://Field/name}

☐ No, my correct NAME is:

Please confirm your MOBILE PHONE/WHATSAPP:

☐ My MOBILE PHONE/WHATSAPP is the one in which I received the link to this form

☐ No, my correct MOBILE PHONE/WHATSAPP is:

(b) *Address Omission*

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

Please confirm your NAME:

☐ My name is {e://Field/name}

☐ No, my correct NAME is:

Please confirm your MOBILE PHONE/WHATSAPP:

☐ My MOBILE PHONE/WHATSAPP is the one in which I received the link to this form

☐ No, my correct MOBILE PHONE/WHATSAPP is:

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

☐ My address from {e://Field/region_embedded} is {e://Field/fullAddress}

☐ No, my address from {e://Field/region_embedded} is:

(c) *Known Address*

Notes: Differences in the application form across the *Status Quo*, *Address Omission* and *Known Address*.


Figure D.3: Job Descriptions

Job Description 1 Sales Representative	Job Description 2 Direct Sales Promoter	Job Description 3 Direct Sales Supervisor
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	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	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

Notes: Job descriptions as presented in the online application forms (translated from Portuguese).

D.4 Information Experiment

Figure D.4: Information Treatment Delivery



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.

D.5 Interview Experiment - 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 their addresses. Hence, when the receptionist 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.6 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. *Is a little careless sometimes*
3. *Is trustworthy*
4. *Tends to be disorganized*
5. *Tends to be lazy*
6. *Perseveres until tasks are completed*
7. *Works efficiently*
8. *Makes and follows 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 professional was the candidate throughout the interview?*

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

- *Had a shaky voice*
- *Stuttered*
- *Laughed nervously*
- *Dressed in informal clothes*
- *Used slang*
- *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 that they knew something 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*

D.7 Self-Assessment Script

The receptionist collected the self-assessments at the reception once the interview was completed. To introduce the self-assessment, the receptionist said:

Hello again, how did everything go? Could you please fill out this form on the tablet so we can learn your thoughts about your interview today?

[receptionist handles tablet to the participant]

[participant reads in the tablet]

We would like you to tell us a bit about how your interview went. Your responses will not be used in your evaluation. They are just for internal records.


[Participants respond on a slider-type of question, from 0-10]

1. Overall, how well do you think you performed at the interview?
2. How nervous were you during the interview?
3. How professional was your behavior during the interview?

Thank you for your answers! Please give the tablet back to the receptionist.

Figure D.5: Spaces Used in the Interview Experiment

(a) Reception



(b) Interview Room



E Theoretical Appendix

Here, we present stylized decision problems that map key elements in each of our experiments. Our goal is not to describe a theory to be tested, but to (i) provide an illustration of how expected discrimination may act in each case and (ii) clarify some mechanisms that can rationalize our results.

E.1 Expected Address Visibility Experiments

Our experiments randomizing expected address visibility should not affect how jobseekers perceive jobs outside our experiment, so we keep the single-job framing presented in Equation 1 in this subsection. In what follows, we borrow the notation introduced in Section 2.2.

E.1.1 Address Omission Experiment

For this illustration, we let the action space $\mathcal{A}_t$ include a binary decision to apply and another simultaneous binary decision about whether to obfuscate one's address (e.g., $o = 0$ favela address; $o = 1$ nonfavela address, with the choice being available only in *Status Quo* applications). We simplify V by assuming it does not depend on o , but it is illustrative to make explicit the possibility that the continuation value of applying to the job depends on whether the jobseeker may suffer address-based discrimination in future stages of the job application procedure, e.g., during an interview. To do that, we will let the expected continuation value of a favela jobseeker who decides to apply be $V(d, d)$ if she is perceived to be a favela resident and $V(0, d)$ if she is not, in a manner that the second d encodes the (possible) consequences of being perceived as a favela resident later on by that employer. We can make some functional form assumptions about the application costs and collapse the continuation probability into the continuation value without loss of generality to write the utility of applying as

$$\mathbb{E}[\pi(o)V(d, d) + (1 - \pi(o))V(0, d) - \alpha(d + o)], \quad (4)$$

with $o \in \{0, 1\}$ being the decision to obfuscate.

We can then think of *Known Address* as a treatment that sets π to $\pi^{KA} = 1$, and *Address Omission* as setting $\pi = \pi^{AO} \leq \pi(1) \leq \pi(0) \leq 1$. Then, if V decreases in its arguments and $\alpha > 0$:

$$\mathbb{E}[U^{KA}(1, \cdot, d)] \leq \mathbb{E}[U^{SQ}(1, o, d)] \leq \mathbb{E}^{AO}[U^{AO}(1, \cdot, d)] \quad (5)$$

because

$$\mathbb{E}[U^{KA}(1, \cdot, d)] = \mathbb{E}[V(d, d) - \alpha d] \leq \quad (6)$$

$$\mathbb{E}[U^{SQ}(1, 0, d)] = \mathbb{E}[\pi(0)V(d, d) + (1 - \pi(0))V(0, d) - \alpha d] \leq \quad (7)$$

$$\mathbb{E}[U^{SQ}(1, o^*, d)] = \mathbb{E}[\pi(o^*)V(d, d) + (1 - \pi(o^*))V(0, d) - \alpha(d + o^*)] \leq \quad (8)$$

$$\mathbb{E}[U^{AO}(1, \cdot, d)] = \mathbb{E}[\pi^{AO}V(d, d) + (1 - \pi^{AO})V(0, d) - \alpha d] \quad (9)$$

In the main text, we mention the possibility that some jobseekers may expect address-based discrimination later on, regardless of their treatment assignments, could explain the average null effects on interview attendance. This would map to perceiving $V(0, d)$ to be very close to $V(d, d)$. We also mention the possibility that nonwhite jobseekers may expect to be recognized as favela residents because of their race, which would map to them perceiving π^{AO} to be close to one.

Also, to the extent that not requiring an address signals that the firm is less likely to discriminate, we might also expect that d is expected to be lower in *Address Omission* than in *Status Quo*. This modification would simply make applying more attractive under *Address Omission*.

E.1.2 Interview Experiment

For the Interview Experiment, we may model the action space as an effort decision and allow for the possibility that it is harder to translate effort into performance when one expects discrimination (e.g., because discrimination induces stress which disrupts one's ability to perform). Hence, a simplified decision problem may look like:

$$\max_e \mathbb{E}[\pi(p(e, d)V(e, d, d) - (1 + d)e) + (1 - \pi)(p(e, 0)V(e, 0, d) - e)] \quad (10)$$

so that effort is more costly when one is perceived as a favela resident and discriminated against.

The mapping of the treatments into the decision problem is similar to the one we proposed for the Address Omission Experiment. Under *Name-Only*, π is small, so the costs of effort are lower. Under *Name-and-Address*, π is large, making it harder to translate effort into performance. These treatments could also have affected candidates' perceived returns to effort, but whether that should increase or decrease effort should depend on the cross-derivatives of the $p(\cdot)V(\cdot)$ terms.

E.2 Information Experiment

Because the information treatments provide signals about callback probabilities across many potential employers (rather than about a single vacancy), we model job search in discrete time with multiple applications per period. Time is $t = 0, 1, 2, \dots$ and jobseekers discount at factor $\beta \in (0, 1)$. Employers are heterogeneous in a discrimination type $d \in [0, \infty)$, distributed according to the jobseeker's subjective belief F . A match with an employer of type d yields value $V(d)$, with $V'(d) < 0$ (more discriminatory employers are less desirable). We abstract from other job attributes.

In each period, an unemployed jobseeker chooses search effort $e \in \{0, 1, 2, \dots\}$, interpreted as the number of applications submitted in that period, and an address choice $o \in \{0, 1\}$, where $o = 0$ means the resume has a favela address and $o = 1$ means a nonfavela address. Let $u(e, o)$ denote flow utility while unemployed, including the cost of submitting e applications. Each application is sent to an employer whose type is an independent draw from F . Conditional on (o, d) , an application is believed to generate a callback with probability $q(o, d) \in [0, 1]$.

Define the implied reduced-form object

$$Q_o(d; F) \equiv \int_0^d q(o, x) dF(x), \quad (11)$$

i.e., the perceived probability that a *single* application results in a callback from an employer with type weakly below d . In particular, $Q_o(\infty; F) = \int_0^\infty q(o, x) dF(x)$ is the perceived *average callback probability per application* under address choice o , for which the results of the audit study are a proxy.

Matching technology. We assume that if at least one callback is received, the jobseeker pursues one callback chosen at random (d is not observable before pursuing the callback). Under independence across applications and draws of employer types, the probability of receiving no callback in the period is

$$\Pr(\text{no callback} \mid e, o) = (1 - Q_o(\infty; F))^e,$$

so the probability of at least one callback is

$$\lambda(e, o; F) \equiv 1 - (1 - Q_o(\infty; F))^e.$$

Conditional on receiving a callback under address choice o , the employer type distribution

is callback-weighted:

$$dF_o^c(d) \equiv \frac{q(o, d)}{Q_o(\infty; F)} dF(d).$$

Hence the expected value of an accepted job conditional on a callback is

$$\tilde{V}_o(F) \equiv \int V(d) dF_o^c(d) = \frac{\int V(d) q(o, d) dF(d)}{Q_o(\infty; F)}.$$

Let U denote the stationary value of unemployment. In a stationary setting, U solves

$$U = \max_{e, o} \left\{ u(e, o) + \beta \left[\lambda(e, o; F) \tilde{V}_o(F) + (1 - \lambda(e, o; F)) U \right] \right\}. \quad (12)$$

Returns to additional applications. For fixed (o, F) , the incremental gain from raising effort from e to $e + 1$ is

$$\Delta_e(o, F) \equiv (u(e + 1, o) - u(e, o)) + \beta (\tilde{V}_o(F) - U) Q_o(\infty; F) (1 - Q_o(\infty; F))^e.$$

The term $Q(1 - Q)^e$ is largest at intermediate perceived callback probabilities and small when success is perceived as either very likely or very unlikely. Hence revisions in perceived callback levels need not imply a single-direction change in optimal e .

To map the information treatments, it is useful to decompose beliefs about callbacks into a baseline component and an address-specific discrimination component. Assume:

$$q(0, d) = \pi \quad \text{and} \quad q(1, d) = \pi + a(d),$$

where $\pi \in (0, 1)$ is a baseline callback level and $a(d) \geq 0$ is weakly increasing with $a(0) = 0$. Then $Q_0(\infty; F) = \pi$, while

$$Q_1(\infty; F) = \pi + \int a(d) dF(d),$$

so the *gap* between nonfavela and favela callback rates identifies beliefs about discrimination (through F and/or the magnitude of $a(\cdot)$) separately from beliefs about the overall callback level π .

Mapping treatments to the model. *Favela Info* provides a signal about the favela callback rate in the audit, which in the model corresponds to $\pi = Q_0(\infty; F)$. Because optimal effort depends on the nonlinear term $Q(1 - Q)^e$, a negative update in π can have a nonmonotonic effect on e .

Full Info additionally reports that the nonfavela callback rate is similar to the favela rate.

Under the decomposition above, this is primarily a signal that $\int a(d) dF(d)$ is small, i.e. that discrimination is less prevalent/severe than jobseekers may have believed. If (i) *Favela Info* locks down beliefs about π and (ii) there is no cross-updating from the nonfavela signal back into beliefs about π , then *Full Info* shifts beliefs about F toward less discriminatory employers (lower d). Since $V'(d) < 0$, this raises $\tilde{V}_o(F)$ and therefore increased the marginal return to effort for a given callback level $Q_o(\infty; F)$ (in particular under $o = 0$, where $Q_0(\infty; F) = \pi$ is pinned down by *Favela Info*). This yields an encouragement effect of *Full Info* on e relative to *Favela Info* absent cross-updating in callback levels.

When cross-updating is present (e.g., jobseekers treat favela and nonfavela callback levels as positively correlated through market tightness), *Favela Info* can also lower beliefs about nonfavela callbacks and *Full Info* can further lower both posteriors. In that case, the net effect of *Full Info* on e need not be signed, consistent with weak behavioral responses despite sizable belief updating.