Making the Elite: Class Discrimination at Multinationals

Soumitra Shukla*

March 25, 2025

Abstract

How salient is hiring discrimination in driving socioeconomic gaps in elite job access, and where do these disparities arise? These questions are particularly challenging for social status discrimination. While widely acknowledged, identifying its source has proven elusive in Western contexts, where social position must be inferred from multiple markers like accent, education, and networks. India's caste system provides a uniquely clear measure of social status, allowing more precise identification of where and how status shapes access to top jobs. Analyzing detailed screening data from over 1000 jobs at US and European multinationals hiring in India, this paper provides the first systematic evidence quantifying both the salience and source of social status or class discrimination in elite hiring. Subjective personal interviews assessing "fit" are the primary mechanism for revealing and penalizing social position, with nearly 90 percent of caste disparities in hiring emerging during these assessments. The higher promotion rates of disadvantaged castes—about 20 percent even in highly selective jobs—suggest that any information gained during personal interviews cannot justify these hiring disparities. These findings reveal how subjective "fit"-focused assessments can systematically exclude high-potential talent based on social status. This presents a fundamental challenge: how can organizations retain necessary managerial discretion while preventing status-based discrimination that undermines both diversity and talent optimization?

*Shukla: Research Affiliate, Department of Economics, Yale University (soumitra.shukla@yale.edu). I would like to thank Joseph Altonji, Costas Meghir, and John Eric Humphries for serving on my committee and providing guidance and feedback. I would also like to thank Jason Abaluck, Alex Albright, Peter Arcidiacono, Crémieux, Michael Dinerstein, Alessandra Fenizia, Arpit Gupta, Patrick Kline, Eliza McNay, Evan Rose, Kamalini Ramadas, Jung Sakong, Rob Valletta, Seth Zimmerman, seminar participants at the Yale labor workshop, the Chicago FED, Chicago Harris, Federal Reserve Board, Wharton (BEPP), Penn State, UCSD, SOLE, Southeastern Micro Labor Workshop, System Equitable Growth Conference at the New York Fed, the "Discrimination in the 21st Century Conference" organized by the Becker Friedman Institute and Chicago Booth, MIT Organizational Economics Lunch, and the NBER Race and Stratification Working Group for helpful feedback and support.

1 Introduction

Who gains access to elite jobs, and how do hiring practices shape inequality in high-stakes recruitment settings? While extensive research documents various forms of hiring discrimination across labor markets (Altonji and Blank, 1999; Rivera, 2012, 2015; Friedman and Laurison, 2020), previous studies lack granular evidence on precisely *where* and *how* discrimination occurs, and its *quantitative salience* in competitive and real-world settings, such as major multinational corporations (MNCs).

This challenge is particularly severe for social status or class (Durlauf and Fafchamps, 2006; Stansbury and Rodriguez, 2024). In Western contexts, social position must be pieced together from multiple, often ambiguous signals—accents that may not clearly map to specific backgrounds, educational credentials that mask family advantages, and social networks that are hard to observe. These overlapping and subtle markers make it difficult to pinpoint exactly how status shapes access to top jobs. Thus, social class discrimination remains comparatively under-studied and more difficult to isolate—especially in high-stakes hiring contexts.

Recent trends in evaluation practices across meritocratic institutions illustrate how well-intentioned efforts can inadvertently reinforce class advantages. For instance, debates around standardized versus holistic assessment reveal this tension. Evidence from Sacerdote et al. (2025) suggests that "test-optional" policies in U.S. college admissions may unintentionally harm high-achieving students from lower socioeconomic backgrounds. Meanwhile, Chetty et al. (2023) finds that many selective private colleges give significant weight to "soft" factors which disproportionately benefit students from high-income families, while not predicting future career success as reliably as traditional academic indicators. Similar dynamics appear in India's elite civil service examinations—widely considered among the most competitive in the world—where subjective "personality" interviews consistently penalize candidates with strong written exam records based on perceived social status, while doing little to predict subsequent job performance (Bhavnani and Lee, 2021).

Although higher-education admissions and public-sector recruitment differ from private-sector hiring, all of these domains rely on a combination of hard metrics and softer, subjective measures. In the same way that dropping standardized tests can obscure academic talent, subjective "fit" conversations at final interviews can veil strong underlying skill—particularly among candidates from lower socioeconomic backgrounds. Although these social class advantages are consistently observed across diverse institutional contexts (Laurison and Friedman, 2019; Rivera, 2012; Stansbury and Rodriguez, 2024), they remain difficult to pinpoint be-

cause measuring "class" precisely is challenging.

To overcome these measurement challenges, this paper leverages recruitment data from leading US and European multinationals in India, where caste provides a uniquely well-defined measure of social status. While affirmative action at top Indian colleges aims to boost the educational attainment of disadvantaged castes, these groups remain largely shut out of elite jobs (Madheswaran, 2008). By pairing granular hiring data with precise measurement of social status, this paper provides the first systematic evidence quantifying the salience and source of status-based discrimination in elite hiring.

The key innovation is documenting how personal interviews—often framed as assessments of candidate "fit"—serve as the primary mechanism for revealing and penalizing social status. By "fit," this paper refers to both legitimate skill- or role-related evaluations and more intangible cultural or background alignment that hiring managers consider relevant during the interview process. Using step-by-step screening decisions, I find that nearly 90% of caste disparities in hiring arise during the final screening round comprising personal interviews. Moreover, once hired, disadvantaged castes have about a 20% higher promotion rate, even in the most selective jobs.

Elite entry-level jobs serve as critical pathways to economic power, making it vital to understand how discrimination shapes access to these positions. In India, this question has particular urgency since nearly 70% of the population belongs to disadvantaged castes. The job fair I study offers a representative window into elite hiring: firms often replicate these practices at other top Indian colleges and in global offices (Bloom et al., 2012).

This setting also provides unique insight into how earlier affirmative action benefits can erode during hiring. Nearly 50% of students in this college where the job fair is conducted are from disadvantaged castes within each major, due to government-mandated affirmative action policies. While recent research shows that improper enforcement of affirmative action can disadvantage protected groups (Yenmez and Sönmez, 2022), this study pinpoints a distinct mechanism: subjective "fit"-focused interviews can reimpose social barriers downstream. By illuminating how subtle discrimination at workforce entry can erode prior gains, these findings underscore that even well-designed affirmative action measures demand attention to institutional structures that can reinstate social barriers.

While caste may be specific to India, the mechanisms underlying discrimination have broader applicability. Caste acts as a uniquely comprehensive marker of social stratification, capturing multiple dimensions of advantage and disadvantage. In line with the concept of economic, social, and cultural capital laid out in Bourdieu (1986)—where economic capital

encompasses financial or material resources, social capital refers to networks and social ties, and cultural capital comprises symbolic assets like education, accent, or mannerisms—caste effectively consolidates these status signals into a single hierarchy.[1] Even when some advantaged caste families are poor (or some disadvantaged caste families are wealthy), caste continues to shape perceived social position (Beteille, 1965, 1969). By contrast, in Western contexts, multiple markers must be pieced together to approximate social class, complicating efforts to identify discrimination.

To understand how interviewers process these status markers, dual-process cognition theories provide valuable insights (e.g., Kahneman, 2011). These theories suggest that decision-making often relies on rapid, intuitive evaluations prone to bias rather than reflective processes. In elite interviews, hiring managers may lean on quick impressions when gauging "fit," thus penalizing candidates whose socio-cultural markers differ from their own. Cognitive scientists highlight how superficial rapport (Todorov, 2017), subtle stereotypes (Knobe and Nichols, 2008), or immediate affective responses (Pinker, 2021) can override more deliberate judgments of a candidate's potential. Our analysis demonstrates these cognitive biases operating within subjective hiring processes.

Conversations about "fit" worldwide often blend intangible cues such as dress style or vocabulary, making it hard to pinpoint which aspects of social identity are being penalized (Rivera, 2015). Caste's clear coding reveals precisely how informal screening criteria translate social position into hiring barriers. As the world becomes more diverse, these mechanisms of status-based exclusion are gaining relevance. Indeed, as Rose (2023) argues, group differences often *endogenously* arise from stereotypes that generate real disparities.

Substantial evidence highlights these patterns: "fit"-focused hiring systematically excludes less privileged groups (Rivera, 2012, 2015, 2020), large-scale experiments uncover systematic employer-level discrimination (Kline et al., 2022), "glass ceiling" effects emerge at higher career rungs (Folke and Rickne, 2016; Macdonald et al., 2024), and in-group biases hamper objective evaluations (Sarsons, 2017; Sarsons et al., 2021).

Even elite institutions struggle to eliminate class-based gaps (Stansbury and Schultz, 2023; Stansbury and Rodriguez, 2024), with career outcomes often determined early (Asher et al., 2024; Durlauf, 1996, 2004; Kaila et al., 2024). Research on Nobel Prize winners shows how privileged access to elite networks perpetuates inequality in academia (Novosad et al., 2024). While these studies demonstrate that social position shapes economic outcomes, they remain hampered by ambiguous class markers in Western contexts.

[1]For further sociological work that expands on Bourdieu (1986) by analyzing how economic, cultural, and social capital shape class-based outcomes, see Lamont and Lareau (1988) and Rivera (2015).

By leveraging a setting where social status is precisely measured, this study delivers two key advances. First, it pinpoints personal interviews as the primary mechanism through which "fit" assessments reveal and exclude candidates based on social status. Second, it provides the first large-scale evidence that these interviews penalize candidates in elite recruitment based on perceived social position.

Each of the 1050 jobs in my sample conducts four screening rounds. The first round involves reading applications. The second comprises written aptitude tests. After testing comes group debates, where students are evaluated on socio-emotional skills. The final stage is the personal interview, where hiring managers screen on potentially many job-relevant skills, including candidate fit (Newman and Thorat, 2012; Rivera, 2012, 2015).

There are large caste disparities among jobs received through the job fair. The caste gap in earnings is 23% without controls and reduces to 10% when controlling for pre-college skills, college performance, and other employer-relevant factors. Critically, because of how the job fairs are set up, there are no caste differences in pay for identical jobs. So the earnings gap is due to differences in job composition across castes. Because salaries offered in the job fair are also *non-negotiable*, the caste penalty in mean earnings at each job search stage can be calculated by comparing the salary differences in the jobs the caste groups are in contention for after each stage.[2] About 90% of the 10% conditional caste earnings penalty occurs between personal interviews and job offers. Similarly, about 95% of the caste penalty in the *success rate* at the job offer stage can be attributed to the personal interview round.

Personal interviews—designed to evaluate extracurriculars, background, and "fit"—are the primary stage where caste disparities emerge. While evaluations at personal interviews may aim to assess job-relevant skills, they also help reveal caste. The clearest evidence that caste itself is being penalized comes from the finding that caste disparities emerge earlier among the small subsample of students whose names plausibly reveal caste.

While some surnames indicate caste, most do not convey caste signals in urban-educated settings—due to migration and large geographical variation in names and naming conventions. The salience of last names in signaling caste is even weaker in elite colleges, as they are typically located in urban centers and draw students from all over India (Jodkha, 2017; Wiltshire, 2020). Consistent with these explanations, only about 15% of students have distinctive caste names. In this small subsample, the caste penalty emerges during the application reading round. However, the *overall* caste penalty during application reading is negligible because most of the sample has non-distinctive caste names. Instead, the penalty

[2]It is common for placement offices in such colleges to enforce a rule preventing negotiation and verify that the offered salaries match those that were initially advertised by firms (Section 3.2).

primarily emerges during personal interviews, where conversations on job-relevant characteristics, including background and cultural fit, help reveal caste. Thus, subjective personal interviews that assess "fit" are the main mechanism through which social status (caste) is revealed and penalized.

The paper's main point is not that personal interviews are useless, but that any extra diagnostic information they provide cannot justify disparities between advantaged and disadvantaged caste interviewees. This is because disadvantaged castes have higher promotion rates once hired. These promotion advantages remain just as large across the job selectivity distribution, where three-quarters of firms hire from only the top 5% of applicants, half from the top 3%, and a quarter from the top 1%. The enduring promotion edge of disadvantaged hires in such highly screened positions strongly indicates that discrimination, rather than unobserved skill deficits, underlies the hiring gap.

To formalize this argument, I develop a simple theoretical framework that leverages job selectivity—how far into the right tail each firm hires. This framework yields two key implications: (1) if disadvantaged candidates truly lacked ability, highly selective jobs should expose that deficit before interviews, but the data reveal no pre-interview caste penalty even in these jobs; and (2) if interviews accurately screened out weaker candidates, then hired disadvantaged workers should not persistently outperform their advantaged peers, yet they do so by about 20% across job selectivity tiers. Using variation in hiring selectivity to link the average promotion rates of hired candidates back to the marginal hiring thresholds in each job, the model demonstrates that these patterns reflect stricter hiring cutoffs for disadvantaged castes—i.e., discrimination—rather than purely skill-based differences.

Given the evidence that employers directly penalize caste, my preferred interpretation is that hiring disparities stem from disparate treatment i.e., from caste-based discrimination. Discrimination based on caste is explicitly illegal under Article 15 of the Indian Constitution (Ministry of Law and Justice, Government of India, 2021).

This paper makes several contributions to the literatures on labor, personnel, and organizational economics by identifying the source of discrimination and quantifying its significance in a high-stakes setting. No other paper thoroughly examines each stage of the entry-level hiring process at leading multinationals, pinpointing both the mechanism through which social status emerges as a barrier (discretionary "fit"-focused conversations) and its magnitude (about 90% of conditional hiring disparities). These findings complement a growing literature on how organizational practices and managerial discretion influence workforce outcomes and productivity (Benson et al., 2019; Bloom et al., 2010; Bloom and Van Reenen, 2007;

Bloom et al., 2012, 2017; Fisman et al., 2018; Friebel et al., 2022; Friebel and Raith, 2024; Hoffman et al., 2018; Hoffman and Tadelis, 2021; Kohlhepp, 2024; Minni, 2023).

This paper advances our understanding of how social status shapes hiring outcomes by leveraging caste as an unambiguous stratification measure. While many studies in sociology have extensively documented how cultural alignment disadvantages less privileged groups (Laurison and Friedman, 2019; Rivera, 2012, 2015), the economics literature on class-based discrimination remains significantly underdeveloped compared to work on race and gender, with only a few notable contributions (Durlauf and Fafchamps, 2006; Durlauf and Ioannides, 2010; Stansbury and Rodriguez, 2024). Recent work by Stansbury and Rodriguez (2024) highlights how even ostensibly meritocratic institutions struggle to remove socioeconomic barriers, yet it does not pinpoint the source of these class-based barriers.

The key advance this study offers is identifying how social status is revealed and penalized through subjective, fit-focused personal interviews. By providing granular evidence in high-stakes elite labor markets, this paper empirically identifies the mechanism through which class discrimination operates. The uniquely well-defined nature of caste as a status marker enables clearer identification of how "fit" assessments both *reveal* and *penalize* social position.

As longstanding markers of social status become increasingly diluted across societies (Clark, 2015), including in India's urban-educated settings, discrimination is likely to push deeper into recruitment processes, making granular data on each hiring stage essential. This paper addresses a key gap in existing research where more diffuse markers would obscure the mechanism through which social status shapes hiring decisions.

The findings have implications beyond "visible" status markers. The way subjective interviews expose and disadvantage candidates based on social background could explain labor market barriers faced by individuals with concealable stigmatized identities, such as LGBTQ+ status (Eames, 2024, 2025).

Discretionary evaluations—particularly personal interviews—have received limited attention in economics, yet this study demonstrates their substantial real-world effects. Hiring practices emphasizing informal screening criteria systematically favor advantaged groups (Hoffman and Stanton, 2024; Rivera, 2015). Given the stakes—actual job attainment rather than initial callbacks—this work significantly advances previous attempts to measure discrimination (Altonji and Blank, 1999). Crucially, it both confirms and refines insights from résumé audit experiments: although early discrimination observed in correspondence studies aligns with that found at the interview stage, anonymization tools may merely shift discrimination

to later phases.

The findings also illuminate important agency problems in hiring. When hiring managers prioritize characteristics like reduced turnover or perceived cultural alignment over true potential, they could generate discriminatory outcomes and organizational inefficiencies. The higher promotion rates of disadvantaged caste hires provide direct evidence that discretionary interviews lead managers to screen out high-potential candidates based on social status, contributing to a broader literature showing how poor hiring practices hinder talent allocation (Gibbons and Roberts, 2013; Lazear, 2009; Lazear and Oyer, 2012; Oyer and Schaefer, 2011).

My results reveal a critical challenge in personnel evaluation: while firms rely heavily on informal screening criteria, these methods could systematically undermine both diversity and talent allocation. Managers commonly use discretionary assessments of fit, yet evidence shows these evaluations correlate poorly with actual performance (Sarsons et al., 2021; Sarsons, 2017; Benson et al., 2023). Current hiring practices often fail to identify top candidates (Cappelli, 2019), but structured, standardized processes can enhance outcomes (Hoffman et al., 2018). Recent work emphasizes both the pivotal role of managerial decisions in shaping firm performance (Fenizia, 2022; Minni, 2023) and how misaligned incentives can exacerbate systemic barriers (Benson et al., 2024; De Janvry et al., 2023; Prendergast, 2007). By documenting how subjective hiring practices simultaneously disadvantage marginalized candidates and undermine accurate talent assessment, this paper demonstrates the broader organizational costs of subjective evaluation methods.

Finally, these findings present opportunities for improving talent allocation and highlight challenges for future research. The evidence that current practices systematically exclude high-potential candidates through subjective assessments suggests scope for organizational innovation in personnel selection. Algorithmic or AI-driven screening tools could inadvertently replicate status-based biases if not carefully validated (Albright, 2024). While India's institutional context provides unique measurement advantages, the tension between discretionary hiring emphasizing "fit" and efficient talent allocation resonates across many settings. Moreover, although Diversity, Equity, and Inclusion (DEI) efforts often focus on race and gender, less than 10% of large US firms track class-based diversity (Stansbury and Rodriguez, 2024), underscoring how class-based disparities remain underexamined. Most directly, the concentration of disparities in subjective interviews suggests that reducing reliance on discretionary screening could enhance efficiency—though balancing standardization with managerial discretion remains an open question.

2 Caste as Social Status Under Affirmative Action

This paper uses two categories of caste—advantaged ("upper") and disadvantaged ("lower")—in line with how India's affirmative action policies classify social disadvantage in public universities and government jobs. Although "caste" might appear distinct from Western notions of "class," the Indian government has long acknowledged the tight interplay between caste, economic marginalization, and social status by grouping together those who have historically faced socioeconomic obstacles under unified administrative categories.

Historical development of caste-based policies. Early steps to uplift socially oppressed groups came via the Government of India Act of 1919, which introduced self-governing institutions and reserved legislative seats for what were then termed "depressed classes," a term meant to capture socioeconomic disadvantage. The Government of India Act of 1935 replaced that term with "Scheduled Castes" (SCs). After independence, India's 1949 Constitution formalized quota-based affirmative action in legislatures, public higher education, and government employment for SCs, Scheduled Tribes (STs), and eventually Other Backward Classes (OBCs).[3]

Caste as a proxy for social disadvantage in affirmative action policies. Although caste began as a birth-based hierarchy, India's policymakers increasingly use these classifications as a measure of "backwardness"—a comprehensive indicator spanning economic, educational, and social disadvantage. In 1979, for example, the Mandal Commission systematically identified socially and educationally "backward" groups and recommended a 27% reservation for them in central and state government jobs, public-sector undertakings, and educational institutions. When these recommendations were implemented in 1990, the overall proportion of constitutionally reserved seats for disadvantaged groups—SCs, STs, and OBCs—rose to nearly 50% (Bayly, 2008).

In effect, the OBC classification explicitly merges caste-based lineage with a broader notion of low socioeconomic standing, demonstrating how caste can function as a consolidated marker of "class" in the Indian context. This design choice—grouping caste with other indicators of disadvantage—underscores that what Western contexts might label "lower social class" is very often encoded as a "backward caste" in India.

Limited scope of affirmative action policies in the private sector. None of these

[3]Scheduled Castes (SCs), also called *Dalits*, and Scheduled Tribes (STs), also called *Adivasis*, are recognized as particularly disadvantaged. Other Backward Classes" (OBC) was an administrative category that the Indian government sought to define more comprehensively following the Mandal Commission's directive to identify the socially and educationally backward classes in India" not already included under SC or ST (Bayly, 2008).

constitutional provisions for uplifting disadvantaged castes, however, applies to private-sector jobs—where the elite US and European multinational positions in this paper are located. Because private-sector roles generally offer higher wages and better career opportunities, and because there are no analogous reservation policies there, disadvantaged castes can remain disproportionately excluded (Subramanian, 2019). This gap between public- and private-sector recruitment highlights how even strong affirmative action at the educational stage can fail to reduce deeper labor-market inequities.

Taken together, these policies reveal how Indian law and popular discourse now treat caste-based disadvantage as broadly analogous to socioeconomic "backwardness." Affirmative action targets individuals historically excluded from economic and educational opportunities, effectively recognizing how caste can reinforce social and economic barriers. Because caste codes social and economic status so definitively, it offers a measurement clarity that Western notions of "class" often lack. By focusing on caste, I therefore capture a rich, historically entrenched measure of "class-like" social status—one that is both legally recognized and practically used to identify disadvantage in India's public institutions.

3 Key Definitions and Institutional Setting

This section clarifies the key terms used in the paper and describes the job fair at the elite college. I conclude by explaining why this job fair provides a representative view of how elite Indian college graduates enter elite jobs.

3.1 Disadvantaged Castes, Elite Jobs, and Elite Colleges

I define "disadvantaged castes," "job," "elite jobs," and "elite colleges" below.

a. **Disadvantaged castes.** Following the Indian government's caste classification, "disadvantaged castes" refer to individuals belonging to Scheduled Castes (SCs), Scheduled Tribes (STs), or Other Backward Castes (OBCs).

b. **Job.** A "job" is a specific job designation within a firm. For instance, if Google hires a product manager and a software engineer, these count as two distinct jobs. Section 4.5 further explains why a "job" typically refers to a firm-designation pair rather than a particular (Indian) office location.

c. **Elite jobs.** An "elite entry-level job" is defined as a formal entry-level position in a multinational corporation (MNC), a major national corporation, or a unicorn (a startup valued at over $1 billion) in the Indian private sector. All jobs in my sample

pay salaries in the top 1% of Indian entry-level wages (Economic Advisory Council to the Prime Minister, 2022). Throughout the paper, I use "elite entry-level jobs" and "elite jobs" interchangeably.

d. **Elite colleges.** "Elite colleges" are those consistently ranked in the top 30 in their respective fields by *India Today*, an Indian equivalent of *US News & World Report*. These fields include science, engineering, arts and humanities, social sciences, and business. Examples include the Indian Institutes of Technology (IITs), Indian Institutes of Management (IIMs), National Institutes of Technology (NITs), and prestigious Delhi University colleges such as St. Stephen's, Shri Ram College of Commerce, and Lady Shri Ram College.

3.2 The On-Campus Job Fair of the Elite College

The elite college's job fair—a key institutional feature of this study—occurs entirely on campus and consists of three phases: (1) pre-placement, (2) placement, and (3) post-placement, also known as the "aftermarket."

a. **Pre-Placement Phase.** The placement office invites firms before June. Between June and mid-August, recruiters visit campus to give pre-placement talks, gauge student interest, and finalize which roles they will offer. By late August, they make return offers to students who interned with them over the summer, known as pre-placement offers (PPOs), which have late-August deadlines. Students who accept PPOs are *disallowed* (de-registered) by the placement office from participating in the full-time job placement process.

How important are return offers from summer internships in this setting? Unlike in the US, return offers from junior-year summer internships play a minor role in full-time job searches at elite Indian colleges. In this setting, fewer than 5% of the graduating cohort accepts such offers and exits the regular on-campus placement process for full-time jobs. Thus, more than 95% of the cohort participates in the job fair, with fewer than 2% choosing public-sector jobs (Section 4.5).

Registration by students participating in the job fair. During the pre-placement phase, students register for the job placement process by late August. By early September, firms submit *employer registration forms* to the placement office, detailing job positions, compensation packages, and projected vacancies for that year. Once submitted, these job profiles are "locked" and cannot be altered. Students are also prohibited from bargaining over compensation packages. The placement office verifies

advertised pay by requiring students to submit copies of their offer letters and can restrict future campus visits by firms that fail to provide the advertised compensation.

b. **Placement Phase.** The placement phase begins in mid-September, when students apply for jobs. Employers conduct initial "first cuts" based on applications, then administer written aptitude tests and group debates (which assess socio-emotional skills) to determine interview eligibility. These three pre-interview rounds—application screening, aptitude testing, and group debates—occur from September to early December. From December to January, employers hold personal interviews, the final screening round, and then make job offers. Students choose among their offers, and by January, about 70% of the cohort secures full-time jobs.

c. **Post-Placement Phase or the "Aftermarket."** From late January until graduation in late July, the placement office runs an "aftermarket" for students who registered for the placement process but did not secure jobs. About 30% of the full sample enters this phase.[4] The aftermarket relies heavily on referrals from the placement office or personal networks. About 80% of these aftermarket participants (25% of the full sample) ultimately pursue entrepreneurship, graduate studies (e.g., Master's or PhD), or gap years to prepare for civil service or elite MBA exams. The remaining 20% of aftermarket participants (about 5% of the full sample) eventually find jobs elsewhere.

What type of firms participate in the aftermarket? Firms hiring in the aftermarket typically differ from those recruiting during the main cycle. They often pay less or are newer startups, and their recruitment processes are less standardized. Because my analysis focuses on the standardized hiring process during the regular placement cycle—where the college collects data and all jobs involve four screening rounds—I exclude the aftermarket and address the implications of this selection in Section 5.2.

3.3 Representativeness of the College's On-Campus Job Fair

The college's job fair provides a representative view of how elite Indian college graduates transition into elite jobs for four main reasons:

First, as defined in Section 3.1, "elite colleges" are those consistently ranked among the top 30 in their respective fields. Nearly all such elite colleges operate similar on-campus job fairs,

[4]Many of these students approach the market "softly," as they may already plan to pursue other opportunities or cast a narrow net during the main hiring cycle. Thus, they often apply only to very selective firms or niche startups.

adopting standardized procedures. Consequently, the rules, protocols, and student-employer interactions at the studied college's fair closely resemble those at other elite institutions.

Second, the same group of elite employers—multinational corporations, leading national firms, and unicorn startups—actively recruits from these top-tier colleges. Firms visit about 30-40 other Indian colleges every year, each with about 1500 students, suggesting that about 50,000 students are impacted annually by the personnel practices of such firms.

Because they source their elite entry-level hires from a similar network of institutions, the job opportunities and recruitment strategies are largely uniform across them. All jobs in my sample are extremely high paying (Section 3.1), ensuring that the studied college's placement outcomes are emblematic of the elite labor market segment as a whole.

Third, elite entry-level jobs almost exclusively hire from these elite colleges, and elite college graduates, in turn, almost exclusively join these elite entry-level positions. This mutual exclusivity—documented in prior research (Subramanian, 2019; Jodhka and Naudet, 2019)—creates a tightly interconnected labor market. The pathways observed at the studied college thus mirror the broader ecosystem of elite transitions in India.

Fourth, only a negligible fraction of elite college graduates forego their college's on-campus job fair. In the studied setting, fewer than 5% of students opt out of the formal placement process for full-time jobs (Section 3.2). Comparable opt-out rates at other elite colleges further confirm that the placement process here closely resembles that of its peers (Subramanian, 2019). Moreover, the vast majority of students rely on the college's placement office rather than off-campus channels, enhancing the fair's representativeness.

In sum, the on-campus job fair at the studied college is not an isolated or atypical phenomenon. Instead, it exemplifies a widely replicated, highly standardized institution that links elite colleges to a common set of top-tier employers.

Online Appendix Section A.1 provides additional details and evidence supporting the representativeness of this setting.

4 Data

This section provides an overview of the data on students, jobs, and job search. I also discuss data on longer-term outcomes such as job promotions and tenure. I conclude by discussing some data limitations.

4.1 Data Source

The data come from the university's placement office, which maintains detailed records of each stage of the job search process, including applications, screening rounds, interviews, and final job offers. This institution qualifies as an "elite college" consistent with the definition in Section 3.2. It is one of India's top-ranked universities, with a placement process widely regarded as a benchmark for other elite colleges. The university's name is kept anonymous to ensure confidentiality.

4.2 Data Completeness

These fairs operate under comparable rules that both firms and students must follow, ensuring a standardized and well-defined hiring process across institutions.

Both firms and students must go through all screening rounds and cannot make or accept offers earlier. Every job in my sample conducts all four screening rounds before extending offers (Section 4.7). The placement office enforces strict protocols: having applied to a job, students must attend every screening round for which they qualify; firms are prohibited from making early or informal offers. These rules, common at other elite Indian colleges, help maintain the integrity of the data. Because all participants must comply, the placement data are accurate and complete. While understanding the origins of these rules is beyond the scope of this paper, they are publicly available on most elite colleges' placement websites.[5]

4.3 Time Span, College Degrees, and Admissions Criterion

Below, we provide a concise overview of the data—covering the four-year time span, illustrating the near-equal caste representation within each college major, and noting the singular role of national-level entrance exams in determining admissions.

Time Span. The data belong to four placement cycles or years. The exact years are pre-pandemic and have been omitted to further preserve anonymity of students and jobs.

Nearly equal representation of both castes within each college major due to affirmative action policies in admissions. Table 1 shows the total number of students belonging to each caste within each college degree. As mentioned in Section 2, students are admitted to public colleges in India through caste-based affirmative action policies. Such

[5]See Central Placement Cell (Delhi University).

policies ensure that the caste share *within each college major* is nearly equal. There are a total of 4164 students in the sample that are about evenly split between castes.

How is caste reported by students? Caste in the data is reported by students, typically by presenting a formal "caste certificate" to the college, and is used for quota-based policies in admissions.

The sample is overwhelmingly male. Male students comprise about 90% of the sample (not shown). This proportion is typical of those at elite technical colleges in India (Datta, 2017).

College degrees. There are four college degrees in the sample (Table 1). They are the Bachelor of Technology (B.Tech.), "Dual," Master of Technology (M.Tech.), and Master of Science (M.S.) degrees. A Dual degree integrates undergraduate and post-graduate studies and is completed a year after the conventional four-year degree (B.Tech.).

College admissions criteria. Students are admitted to the elite college *solely through national-level entrance exams* based on caste-major-specific cutoffs.

Table 1: **Total Students by Caste and Degree**

	Total Students		
	Advantaged Caste	Disadvantaged Caste	Degree Total
B.Tech.	579	710	1289
Dual Degree	622	617	1239
M.Tech.	601	566	1167
M.S.	344	125	469
Caste Total	2146	2018	4164
Fraction	0.52	0.48	1.00

Notes: Table 1 shows the total number of students by caste within each college degree, the total number of students by caste, and the fraction of students by caste. B.Tech. stands for Bachelor of Technology, M.Tech. stands for Master of Technology, and M.S. stands for Master of Science. A Dual degree integrates undergraduate and post-graduate studies and is completed a year after the conventional four-year degree (B.Tech.).

4.4 Overview of Student Characteristics

Below, we describe three key student characteristics—college GPA, pre-college skills, and previous labor market experience—and then highlight two main patterns that emerge from the data.

1. **College GPA.** The data includes the college GPA of students. There are substantial caste differences in college GPA within each degree (Table 2). Importantly, there are

students from both castes (i.e., common support) within each GPA decile (Online Appendix Figure OA.1).

2. **Pre-college skills.** The data includes entrance exam (EE) scores of students. EE scores are the only criterion for college admission (Section 4.3). EE scores in the data are *originally exam ranks* that have been renormalized so that larger numbers are better. Caste differences in EE scores are large (Table 2). As with college GPA, there are students from both castes within each EE score decile (not shown).

3. **Previous labor market experience.** The data also include detailed information on both summer and winter internships, including duration of internship employment, duration of part-time or full-time employment, total pay during internships, total pay during part-time or full-time employment, sector of internship employment, and employment in startups. Internship descriptions typically include application eligibility criterion and desired skills. Application eligibility depends upon a combination of factors such as major, degree, and GPA.

There are two main patterns that emerge from the data on student characteristics:

1. **Large caste differences in college GPA and college entrance exam scores.** Table 2 shows that there are large caste differences in college GPA but students from both castes within each GPA decile (Online Appendix Figure OA.1). Similar results hold for college entrance exam scores (not shown).

2. **Weak correlation between caste and measures of previous labor market experience.** For example, consider a key measure of previous labor market experience: sector of internship employment. There are modest caste differences in the sector of internship employment among both B.Tech. and Dual Degree students. Regardless of caste, about 25% of these students interned in the technology sector, 37% in consulting, and 38% in manufacturing. Similar results hold for students pursuing M.Tech. and M.S. degrees (not shown).

   The weak correlation between internship outcomes and pre-college skills is a common pattern at elite Indian colleges. Because return offers from junior year summer internships are not an important feature of full-time job search at elite Indian colleges (Section 3.2), students seem to care more about the job sector or the firm instead of salaries; apart from a select few firms, students are just happy to get any internship and vice-versa. This could partly explain modest caste differences in internship outcomes.

Table 2: **Summary Statistics (Students)**

	Adv. Caste	Disadv. Caste	Diff (SD)
		B.Tech.	
Avg. Entrance Exam Score	0.41	-0.33	0.74***
Avg. 10th Grade Score	0.24	-0.19	0.43***
Avg. 12th Grade Score	0.24	-0.20	0.44***
Avg. Overall College GPA	0.51	-0.42	0.93***
		Dual Degree	
Avg. Entrance Exam Score	0.35	-0.35	0.70***
Avg. 10th Grade Score	0.20	-0.20	0.40***
Avg. 12th Grade Score	0.14	-0.14	0.28***
Avg. Overall College GPA	0.43	-0.43	0.86***
		M.Tech.	
Avg. Entrance Exam Score	0.27	-0.30	0.57***
Avg. 10th Grade Score	0.23	-0.15	0.38***
Avg. 12th Grade Score	0.17	-0.10	0.27***
Avg. Overall College GPA	0.34	-0.33	0.67***
		M.S.	
Avg. Entrance Exam Score	-0.02	0.02	-0.04
Avg. 10th Grade Score	0.03	-0.07	0.10
Avg. 12th Grade Score	0.02	-0.06	0.08
Avg. Overall College GPA	0.04	-0.14	0.18*

Notes: Table 2 documents some select summary statistics for students on measures of pre-college skills and college academic performance. Pre-college skill measures include college entrance exam scores and scores in the 10th and 12th standard national-level exit exams. College academic performance is measured by college GPA. All scores are pooled and normalized to have zero mean and unit standard deviation within each degree and year. College entrance exam scores are originally ranks that have been renormalized so that higher numbers are better. The difference across castes is reported in standard deviation units. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$. Adv. Caste stands for advantaged caste and Disadv. Caste stands for disadvantaged caste.

4.5 Overview of Job Characteristics

There are about 1050 jobs in the sample, with only about 2% in the public sector and the rest in the private sector. Most jobs come from major multinational corporations—often headquartered in the US or Europe and commonly featured in the Fortune 500—but recruit for positions in their Indian offices. Table 3 provides a snapshot of these jobs. Over the four placement cycles in the data (Section 4.3), a given job may not recruit every year: some do so only once, while others come back during all four cycles. On average, each job appears in about two out of the four cycles. Around 75% of the positions are non-client-facing roles (e.g., software or engineering), whereas the remaining 25% are client-facing (e.g., consulting or managerial)

Omitting public-sector jobs. I omit public-sector jobs from my analysis for two reasons. First, such jobs comprise less than 2% of all jobs in the sample (Table 3). Second, public-sector jobs in India are quite different from their private-sector counterparts, especially in areas like salary structure and job stability.[6]

Salary posting and national wage setting. Jobs post uniform, job-specific (not match-specific) salaries that are non-negotiable during the course of job search (Section 3.2). Conditional on college degree, job salaries do *not* vary across major, caste, or gender (or any other candidate-specific factors).

In fact, firms appear to set salaries *nationally*, perhaps due to administrative convenience. The same firm will be touring multiple colleges for the same job profile and standardization of the compensation package for that job profile across all its Indian locations ("national wage setting") is likely a convenience measure.[7]

Job descriptions and non-pecuniary amenities. Employers submit detailed "employer registration forms" before recruitment begins, outlining everything from job responsibilities and skill requirements to non-pecuniary amenities like health benefits, travel allowances, housing assistance, or signing bonuses. These non-wage attributes are generally standardized within each advertised position—just like the salary component.

Job location. Over 98% of the jobs in my sample are in the Indian offices of firms. When hiring for Indian job positions, employers conduct screening rounds for all Indian offices at once. The specific Indian office that a candidate ultimately gets assigned to is decided later

[6] See the report of the Seventh Central Pay Commission, 2016.

[7] More details about national wage setting in this context are available upon request. National wage setting has also been found among elite entry-level jobs in the United States (Bloom et al., 2012; Sarsons et al., 2022).

and usually depends on the location-specific needs of that employer. For example, Google will conduct a combined recruitment process for its software engineering role regardless of whether the candidate may ultimately get assigned to either the Mumbai or Bangalore office of Google. Therefore, a "job" will almost always mean a firm and job designation pair and not the ultimate (Indian) office location. Moreover, as mentioned above, the compensation package for a given employer and job designation pair is typically standardized across all its Indian locations. The standardization of compensation packages is also compatible with how recruitment is conducted for an employer's Indian locations.

For the firms in my sample that hire for their non-Indian locations (less than 2%), I consider them as a separate job. When hiring for non-Indian locations, an employer will specify an office location. Recruitment processes for firms hiring for non-Indian locations are typically conducted by recruitment teams different from the ones that hire for the same firm's Indian locations; however, such recruitment teams still mostly comprise Indian hiring professionals. Compensation packages of firms that hire for non-Indian locations also differ from those offered by them for their Indian locations. For example, Google hiring a software engineer for its Palo Alto office would offer a different compensation package than if it were hiring for the same position for its Mumbai office. Therefore, (Google, software engineer, Palo Alto) is a different job than (Google, software engineer, Mumbai). However, given the standardization of both the recruitment process and the compensation package across an employer's Indian locations, it is not meaningful to distinguish between (Google, software engineer, Mumbai) and (Google, software engineer, Bangalore).

Hiring managers are typically distinct from direct supervisors. Recruitment, including personal interviews, is typically conducted by hiring managers rather than direct supervisors. The average recruiter travels across multiple colleges and geographic areas to conduct hiring (Section 3.3), making it impractical for direct managers to be heavily involved. Informal conversations with employers recruiting from this college confirm that hiring is primarily managed by dedicated recruitment professionals or managers rather than those who directly oversee the new hires on the job.

Most hiring managers are Indians. Since more than 98% of the jobs are located in India, it is reasonable to assume that recruitment is typically conducted by Indian hiring professionals or managers.

Return offers from internships. Recall, these offers have late August deadlines and students who accept them are disallowed ("de-registered") by the placement office from participating in the formal placement process for full-time jobs (Section 3.2). Therefore,

a typical student who applies to a full-time job during the regular placement season that begins in September would *not* have completed a summer internship in his junior year at the same job.

Table 3: **Summary Statistics (Jobs)**

	Count	Fraction	Avg. Salary ($)	Std. Dev. ($)
Private	1011	0.98	60341.38	32591.51
Government	19	0.02	56202.23	20240.01
Headquartered Outside India	620	0.60	64562.74	36038.04
Headquartered in India	410	0.40	53766.04	23946.16
Client Facing	246	0.24	57553.23	28288.23
Non-Client Facing	784	0.76	61115.92	33361.87
All Jobs	1030	1.00	60265.02	32391.89

Notes: Table 3 shows the proportion of jobs in the public versus private sector, the proportion of jobs headquartered outside versus in India, and the average salary breakdown (USD PPP) among client-facing jobs and non-client-facing jobs. Firms headquartered outside India are mostly based in the US or Europe.

4.6 Data on Job Selectivity and Screening Rounds

There is substantial variation across jobs in the cumulative proportion of students cut both at as well as by the end of each screening round. Table 4 shows the distribution of both the conditional and cumulative cuts across jobs. A "conditional cut" is the proportion of candidates cut by a given job *at* a particular screening round. A "cumulative cut" is the total proportion of candidates cut by a given job *by the end of* a particular screening round. For example, if a job cuts 50% of the candidates at each of the first two screening rounds, then the conditional cut made by the job at each round is 50%, whereas the cumulative cut made by the end of the second round is 75%. Students almost always have to be within the top 5% of a given job's initial applicant pool to get offers. Note that a decent proportion of jobs often do not make any offers, as the job search sometimes fails for many reasons.

Table 4: **Variation in Conditional and Cumulative Selectivity Across Jobs**

Cut Type	Round	Q1	Mean	Q2	Q3
Conditional Cut	App. Reading	0.30	0.39	0.39	0.48
Cumulative Cut	App. Reading	0.30	0.39	0.39	0.48
Conditional Cut	Written Test	0.10	0.18	0.15	0.26
Cumulative Cut	Written Test	0.38	0.49	0.49	0.61
Conditional Cut	Group Debate	0.69	0.74	0.74	0.79
Cumulative Cut	Group Debate	0.83	0.87	0.87	0.91
Conditional Cut	Personal Interview	0.61	0.75	0.78	0.90
Cumulative Cut	Personal Interview	0.95	0.96	0.97	0.99

Notes: Table 4 shows the distribution of both the conditional and cumulative cuts across jobs.

4.7 Details on Screening Rounds

The data comprise the list of job applications, written aptitude tests, group debates, interviews, job offers, and final job choices for each student. The job search typically begins in September when students submit applications (Section 3.2).

Job applications are effectively costless during to a centralized application portal. Students apply for jobs through a centralized job application portal, similar to Job Openings for Economists. Applying to a job only involves "clicking" on the name of the job in the online job application portal and does not require additional cover letters or other statements. Employers only request student resumes that are automatically made available to them when the student clicks on the name of a job to apply. The streamlined job application process significantly reduces the marginal cost of an application. Combined with equal caste representation within each major (Section 4.3), this implies that eligible students from different caste groups apply to jobs at similar rates.

Application eligibility. This depends upon a combination of factors such as major, degree, and GPA. For example, a software engineering position at Microsoft will typically restrict applications to students from the Computer Science major. On the other hand, an analyst position at Goldman Sachs will typically invite applications from all college majors.

Such application eligibility criteria are typically why, despite the low marginal cost of an application, a student cannot apply to all jobs recruiting from campus in a given year.

Number of screening rounds per job. All jobs in my sample conduct four screening rounds and eliminate students at each of them. Each job conducts the following four screening rounds:

1. **Application reading.** In the first round, employers read applications and determine whether to advance students further.
2. **Written tests.** The second round comprises written aptitude tests that are typically conducted online.
3. **Group debates.** After testing comes group debates. During such debates, groups of thirty to thirty-five students are divided into teams of fifteen to twenty people and are made to discuss a shared, non-technical prompt, with each team arguing one side—i.e., a standardized assignment. These debates evaluate applicants on a broad range of socio-emotional skills, such as communication, collegiality, mannerisms, confidence, and consensus building. Written tests and group debates are conducted from late September to late November.
4. **Personal interviews.** After these three increasingly subjective screening stages comes the final screening stage: the personal interview. Such interviews assess candidate "fit" and potentially screen on many job-relevant skills. They are usually conducted between December and January.

Firms typically conduct common screening rounds for all job profiles. For example, if Nvidia advertises 5 job profiles in a given year, it might choose to conduct a common written test ("technical") for all profiles at once.

See Online Appendix Section A.2 for more details on each screening round.

4.8 Summary Statistics on Candidate Screening

Table 5 shows that the average (median) number of applications per job is 208.85 (140.17). Similarly, the average (median) number of written tests conducted, group debates conducted, interviews conducted, and job offers extended per job are 124.91 (87.00), 101.54 (70.00), 26.07 (18.00), and 6.49 (4.00), respectively. Jobs typically restrict applications based on major, degree, and GPA requirements, which limits the number of positions each student can apply to (Section 4.7).

Table 5: **Summary Statistics (Candidate Screening)**

Round	Mean	Median
Application Reading	208.85	140.17
Written Tests	124.91	87.00
Group Debates	101.54	70.00
Personal Interviews	26.07	18.00
Offers	6.49	4.00

Notes: Table 5 shows the mean and median number of applications, written tests conducted, group debates conducted, interviews conducted, and offers extended per job. For a given screening round, these numbers can be calculated by averaging, per job, the number of students across years and then averaging across jobs.

4.9 Data on Job Tenure and Promotions

Linking the sample to data from an employment-focused online platform, I measure two post-hiring indicators of job performance: tenure and promotion rates.[8] I match over 98% of students who received jobs through the college job fair using their names, cross-referenced with a unique combination of internship details, first job, major, degree, and graduation year. Tenure is measured in months as the total duration spent at a firm, calculated from reported start and end dates for each position.

Classifying promotions. Promotions are defined as movements to higher job designations within the same firm, where I classify job titles as vertical advancements based on two criteria. First, I cross-reference job titles with public salary data from platforms like Glassdoor to confirm that the new designation is associated with higher pay. Second, I verify if the job transition aligns with standard career progression timelines for the given role. Structured pay systems in high-skilled, entry-level jobs at elite multinational corporations (MNCs) make this classification particularly reliable, as promotions within these firms are globally standardized and tied to rigid salary structures (Bloom et al., 2012; Minni, 2024). For instance, in a typical consulting firm, moving from "Analyst" to "Associate" within two years would constitute a promotion based on both salary and hierarchy.

Capturing vertical instead of lateral moves. This approach ensures that transitions classified as promotions reflect vertical career progression rather than lateral moves. These

[8]Name of the platform anonymized for confidentiality.

measures, combined with structured pay practices in elite MNCs, minimize subjectivity and provide robust measures of job performance.

4.10 Data Limitations

There are a few data limitations. First, I do not have access to student resumes. However, given that these resumes follow a standardized format mandated by the placement office, they likely contain information already available to me. Second, I do not observe performance measures or scores at any screening round. I address these limitations in Section 5.6.

5 Penalizing Social Status Through "Fit" Assessments

This paper documents five key findings on how personal interviews reveal and penalize social status (caste). First, a sizeable—yet likely conservative—earnings gap exists across castes. Second, about 90% of the hiring disparities emerge only after "fit"-focused personal interviews. Third, job-level caste penalties confirm that these disparities consistently arise during interviews, across nearly all positions. Fourth, employers penalize caste once it becomes apparent, making personal interviews the central channel through which social status is revealed and penalized. Finally, disadvantaged castes achieve higher promotion rates, even at the most selective jobs hiring only the top 1–2% of applicants—implying that discrimination, rather than unobserved ability, underlies the observed hiring inequalities.

5.1 Main Sample

As noted in Section 3.2, roughly 70% of students secure full-time jobs by early January, with fewer than 5% opting out via summer-internship offers. Consequently, over 95% of the cohort participates in the on-campus job fair, resulting in about 2,550 hires (65% of the entire sample) through a standardized, four-round screening.

Initial job placements are reliable measures of post-graduation employment. Since nearly all hires begin the roles they obtained during placement, the placement office's data reliably capture actual post-college employment.[9]

[9] A mandatory exit survey (98% participation) conducted two months post-graduation indicates that hires started in their offered roles without renegotiating compensation. Further checks via a leading employment-focused platform corroborate these findings, ensuring the accuracy of the placement office's records.

5.2 Large Earnings Gap Across Castes

Disadvantaged castes earn less: the unconditional earnings gap is 23%, which decreases to 10% once I control for pre-college skills, academic performance, prior labor market experience, and other employer-relevant traits. This 10% gap—approximately $6,000 (PPP)—reflects differences in job composition rather than unequal pay within identical positions, because the job fair sets uniform compensation across the same roles (Section 4.5). I also find modest variation by job type (Online Appendix Table OA.1).

Discussion of students omitted from the earnings regression. The omitted students fall into two groups: those accepting summer internship offers and those who remain unplaced in the regular cycle and move to the aftermarket (Section 3.2).

Omitted students and selection effects. No registered student rejects all offers. However, as mentioned in Section 3.2, two groups are excluded from the main sample: (i) those accepting summer-internship offers, and (ii) those remaining unplaced after the regular cycle and entering the "aftermarket" (about 30% of full sample). Among the omitted sample, Online Appendix Table OA.2 shows disadvantaged castes are more negatively selected on GPA than their advantaged peers.[10] Since higher GPA typically predicts higher pay (not shown), omitting these disproportionately lower-GPA disadvantaged students understates the true wage gap.

5.3 Caste Earnings Penalty Is Due to Personal Interviews

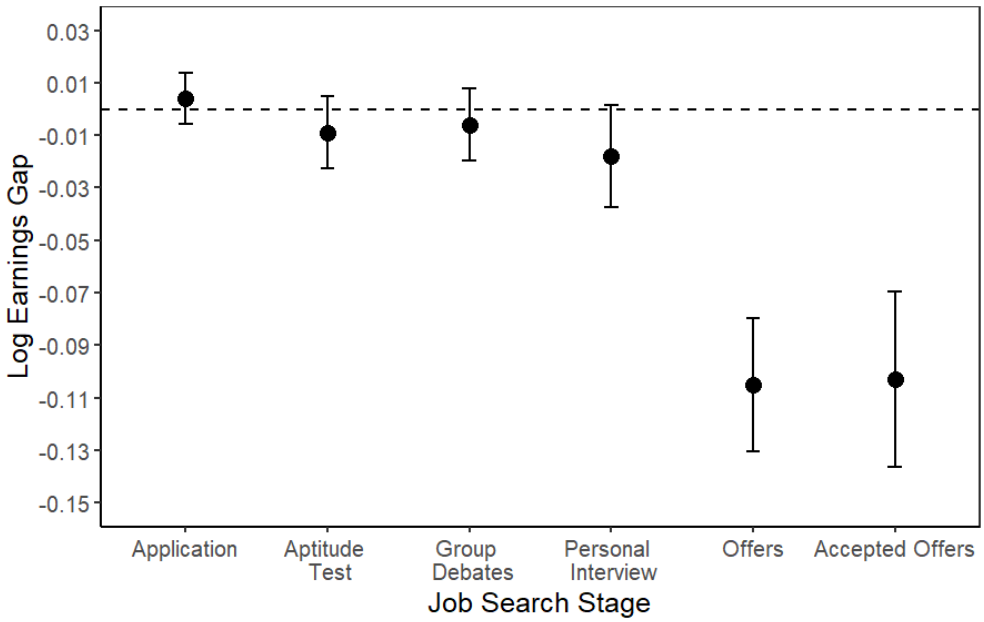
Sections 5.3–5.6 detail the key innovation of this paper: personal interviews, especially their emphasis on subjective "fit" (communication style, personality, background, etc.) serve as the primary mechanism for caste- or status-based discrimination in elite hiring, accounting for nearly 90% of the conditional hiring disparities.

Since each position offers a fixed, non-negotiable salary (Section 4.5), I can decompose the caste earnings penalty by comparing the average posted salaries of roles that remain in contention for each group at each screening stage.

This analysis shows that nearly 90% of the 10% conditional earnings gap emerges after personal interviews, pinpointing late-stage "fit" assessments as the main hiring stage where caste hiring disparities emerge.

[10] *Within* any job offered through the placement process, however, disadvantaged hires tend to have slightly higher GPAs than advantaged hires (Online Appendix Table OA.3).

Figure 1: **The Caste Earnings Penalty**



Notes: Figure 1 shows the log earnings gap across castes at successive stages of job search. Each coefficient in the figure is represented by a black dot and reports the percentage difference in the average salary at each job search stage between disadvantaged and advantaged castes. The vertical bars are 95% confidence intervals. These regressions include controls. There are no caste differences in pay for identical jobs and since offered salaries are non-negotiable during the job fair, the caste penalty in mean earnings at each stage can be calculated by comparing the salary differences in the jobs the caste groups are left in contention for after each stage.

To illustrate this, I estimate variations of the following regression:

$$\log(\text{Avg. Job Salary}_{it}^{\text{Search Stage}}) = \alpha + \beta \times \text{Disadv. Caste}_i + \text{Controls}_i + \gamma_t + \epsilon_{it}, \quad (1)$$

where $i \in \mathcal{I}$ indexes students, γ_t denotes year fixed effects, and Search Stage $\in$ {Application, Aptitude Tests, Group Debates (GD), Personal Interviews, Offers, Accepted Offers}. The controls are characteristics such as college GPA, college entrance exam score, previous labor market experience, major, and degree. For each student i, "average job salary" at a given job search stage s is calculated by averaging over the posted salaries of jobs j that she remains in contention for at that stage.

The coefficient of interest is β, which is shown in Figure 1 for each job search stage. Online Appendix Table OA.4 reports these regressions with Oster (2019) bounding checks that

account for unobserved confounders.[11]

Robustness. While these findings document the overall caste earnings penalty, many factors could confound the average analysis. To address this, Section 5.6 exploits *job-level heterogeneity* in selectivity and shows that disadvantaged castes exhibit no early-stage gap ((application, written tests, group debates) in success rates—even in the most selective positions. This result rules out a "right-tail deficit" explanation prior to personal interviews.

In the next subsection, I show that these late-stage penalties also arise in success rates, reinforcing the primacy of personal interviews in driving caste disparities.

5.4 Caste Penalty in the Success Rate Is Due to Interviews

While the earnings penalty analysis highlights when disparities emerge, success rates provide a complementary lens that accounts for job-level heterogeneity. The findings in Section 5.3 document that negligible earnings penalty exist until personal interviews, based on the composition of jobs that each caste group remains in contention for. Given standardized, non-negotiable salaries throughout the placement process, we should logically observe no success rate penalties *within* jobs prior to interviews. If both caste groups face similar potential salary distributions at stage $s+1$, they must have experienced similar within-job success rates at stage s — otherwise, we would observe differences in the composition of accessible jobs and thus differences in potential earnings. The following analysis confirms this prediction by examining caste differences in success rates at each screening round and finding no difference before interviews.

Using success rates as an alternative outcome to earnings. We define the "success rate" as the probability that student i passes screening round s for job j in year t. As shown in Table 6, more than 95% of the caste penalty in the final job-offer success rate can be traced to the personal interview stage, reinforcing that interviews are the key juncture where disadvantaged groups incur hiring penalties.

To demonstrate this, I run standard linear probability regressions with job fixed effects. Table 6 reports these results, along with bias-adjusted coefficients using Oster's method. Each observation corresponds to a student who remains in contention at a particular screening

[11] We apply Oster (2019)'s bounding approach to address potential omitted-variable bias. Following Oster's recommended benchmarks, we set $\delta = 1$ (i.e., unobservables are assumed to have the same degree of selection as observables) and $R^* = 1.3 \times \tilde{R}$, allowing for the total explanatory power to increase by up to 30% beyond the observed R^2. Under these assumptions, the bounds remain consistent with our main results.

round for a given job and year. For instance, the first column includes all applicants, the second all written test takers, and so forth. Specifically, I estimate the following linear probability model:

$$\Pr(\text{Success}_{ijt}^{\text{Job Search Stage}}) = \alpha + \beta \times \text{Disadv. Caste}_i + \text{Controls}_i + k_j + \gamma_t + \epsilon_{ijt}, \qquad (2)$$

where $i \in \mathcal{I}$ indexes students, $j \in \mathcal{J}$ indexes jobs, Search Stage $\in$ {Application, Aptitude Tests, Group Debates (GD), Personal Interviews, Offers, Accepted Offers}, γ_t denotes year fixed effects, and k_j denotes job fixed effects.[12] The controls are characteristics such as college GPA, college entrance exam score, previous labor market experience, major, and degree.

The results confirm that caste differences in success rates stay negligible at the application, written-test, and group-debate stages, only to reappear at personal interviews. As shown in Online Appendix Tables OA.5–OA.6, these patterns hold for both client-facing and non-client-facing jobs.

Robustness. A natural concern is that any early-stage penalty might be masked by disadvantaged castes having lower ability only in a narrow, extreme tail. However, Section 5.6 confirms that even in the most selective jobs—those with the strictest early-round cutoffs—no caste gap emerges until personal interviews, where caste finally becomes apparent. This pattern underscores that the final stage is the key juncture at which the gap arises, rather than any unobserved skill deficit prior to personal interviews.

Having established that personal interviews coincide with both the earnings and success-rate penalties, I next investigate *when* and *how* caste is revealed at interviews.

[12] See Section 4.5 for how a "job" is defined.

Table 6: **Caste Differences in the Success Rate at Each Job Search Stage**

	Success App (1)	Success Test (2)	Success Debate (3)	Success Interview (4)
Disadv. Caste (OLS)	-0.0025*	-0.0002	-0.0005	-0.2938***
	(0.0014)	(0.0018)	(0.0027)	(0.0048)
Disadv. Caste (Oster)	0.009	0.010***	-0.002	-0.296***
	(0.0013)	(0.002)	(0.003)	(0.005)
Mean Success Rate	0.5990	0.8110	0.2580	0.2610
Observations	305,541	183,150	148,489	38,378
Adjusted R^2	0.61069	0.34360	0.01858	0.24450
Year FE	✓	✓	✓	✓
Job FE	✓	✓	✓	✓
Other Controls	✓	✓	✓	✓

Notes: Table 6 shows the caste difference in the success rate at each stage of job search. Each observation is a student who remains in contention at a given screening round in a given job in a given year. Each column reports a separate linear probability regression and includes controls. These regressions also include job fixed effects. Note that some firms conduct the same screening round (written tests, group debates, interviews) for multiple job positions. However, since these job positions within the same firm represent different jobs, the regressions consider these jobs and the corresponding advance-or-out decisions made by them separately. Bias-adjusted coefficients are calculated using Oster (2019). Following Oster's recommended benchmarks, I set $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$, implying that unobservables are assumed to have the same degree of selection as observables, and the total explanatory power can increase by up to 30% beyond the observed R^2. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

5.5 Interviews as the Main Mechanism for Revealing and Penalizing Social Status

The results documented in Sections 5.3 and 5.4 beg an obvious question: when does caste get revealed to employers? I show that subjective "fit" interviews are the main mechanism through which social status (caste) is revealed and penalized.

5.5.1 What Does the Employer See At Each Screening Round?

I first clarify what the employer can reasonably see about the candidate at each screening round.

1. **Application Reading.** Employers review standardized CVs containing the candidate's name, college GPA, work experience, and extracurricular activities. Caste is not declared, and college entrance exam scores are not provided.

2. **Written Tests.** Employers gain online test scores reflecting technical skills required for the job, but no further personal details.

3. **Group Debates.** Employers observe—but do not extensively interact with candidates—after an initial setup. They observe accents, appearance, dialects, and demeanor. However, these cues are weak indicators of caste in urban-educated settings, such as elite colleges (Section 5.5.3).

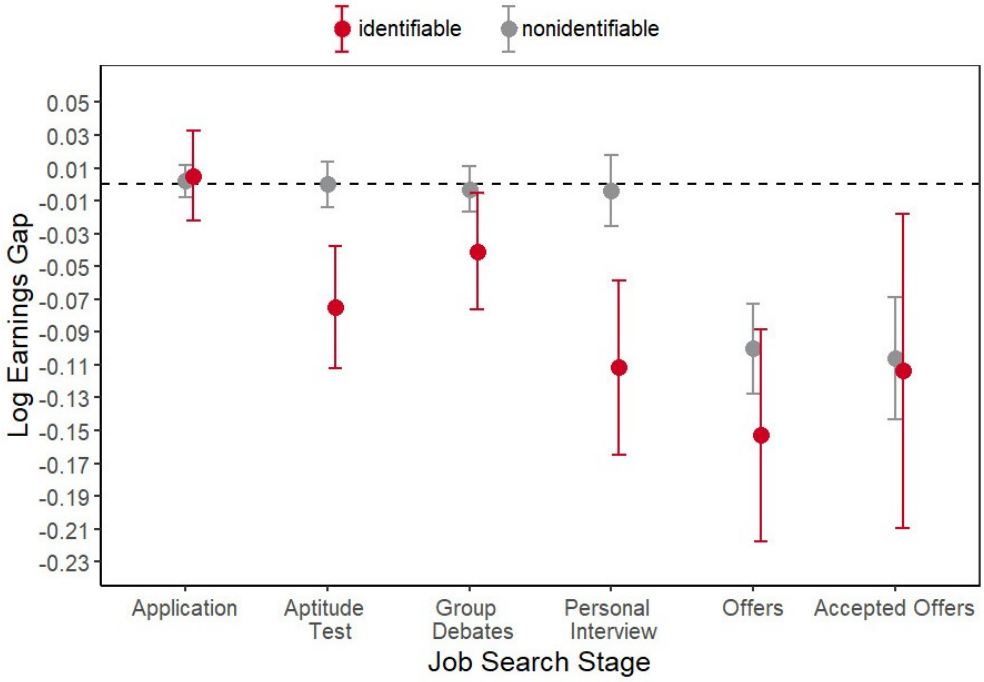
4. **Personal Interviews.** These interviews represent the only extensive one-on-one interaction between the employer and the candidate. Such conversations screen students on potentially many job-relevant characteristics, including candidate background, extracurriculars and fit, that are strongly correlated with caste.

5.5.2 Quantitative Evidence

I provide quantitative evidence linking the caste penalty to the stage at which caste becomes salient. While many surnames are indicative of caste, most do not reliably convey it in urban, elite college settings, where students hail from diverse regions and naming traditions (Jodkha, 2017; Wiltshire, 2020). Consequently, only about 15% of students in my sample have surnames that clearly identify their caste. In that small subgroup, the penalty emerges earlier, during application reading. For the remaining 85% with non-identifiable surnames, the overall caste penalty at the application stage is negligible; instead, it primarily arises during personal interviews, where students are potentially screened through conversations on many job-relevant characteristics, including candidate background and cultural fit, that also help reveal caste. Thus, personal interviews—through informal "fit" conversations—are the main mechanism for uncovering and penalizing social status.

Caste earnings penalty emerges earlier when names reveal caste. Figure 2 replicates Figure 1, dividing the sample into students with clearly identifiable caste names versus those without. I re-estimate Equation 1 separately for each group and plot the coefficients with 95% confidence intervals. The results for non-identifiable names (gray) confirm that disparities do not appear until the personal interview stage, whereas among identifiable names (red), the penalty arises as early as application reading. Online Appendix Tables OA.7 and OA.8 present the corresponding regressions, including bias-adjusted coefficients using Oster's method.

Figure 2: **The Caste Earnings Penalty by Subsample**



Notes: Figure 2 shows the log earnings gap across castes at successive stages of job search in the identifiable (red) versus the non-identifiable (gray) sample. Each coefficient in the figure is represented by a black dot and reports the percentage difference in the average salary at each job search stage between disadvantaged and advantaged castes. The vertical bars are 95% confidence intervals. These regressions include controls. There are no caste differences in pay for identical jobs and since offered salaries are non-negotiable during the job fair, the caste penalty in mean earnings at each stage can be calculated by comparing the salary differences in the jobs the caste groups are left in contention for after each stage.

Same pattern with success rates. The early penalty for identifiable names also appears in success rates. As shown in Online Appendix Figure OA.2, students with identifiable caste names incur a substantial penalty immediately at application reading—across nearly all jobs—whereas those with non-identifiable names face no such penalty.

How are student names associated with caste? To identify clearly caste-linked surnames, I rely on multiple official sources. For advantaged castes, I consult surname lists compiled by the Anthropological Survey of India, while for SC/ST/OBC groups, I draw on (i) Election Commission records from reserved parliamentary constituencies designated for these castes and (ii) official lists of SC/ST/OBC communities maintained by Indian states.

5.5.3 Qualitative evidence supporting caste revelation at interviews

Supporting our quantitative findings that caste penalties emerge primarily at the interview stage, qualitative evidence further indicates that personal interviews, rather than group de-

bates or earlier rounds, enable employers to discern caste. During group debates, several factors limit caste identification: these involve large mixed groups discussing non-technical topics with minimal employer interaction; every participant is assigned the *same standard task* rather than receiving personalized questions; research shows no strong correlation between physical features and caste among urban educated elites (Mishra, 2015; Parameswaran and Cardoza, 2015); and the English language used in these interactions lacks caste-specific markers (Wiltshire, 2020). Even for attentive employers, accurate caste assessment would be practically impossible during these group settings, as visual, linguistic, and nominal indicators offer minimal reliable information about social position in these elite, urban contexts.

Unlike large-group interactions with limited personal detail, interviews typically provide employers the opportunity to delve into a candidate's background, preferences, and cultural fit—factors strongly linked to social position in elite, urban India. Personal interviews routinely explore candidates' biographical details, recreational activities, urban lifestyle preferences, travel experiences, and other elements of cultural compatibility. Research consistently demonstrates these social and cultural markers strongly correlate with social position or caste in India's elite educational environments (Jodhka and Newman, 2007; Deshpande, 2011; Jodkha, 2017; Jodhka and Naudet, 2019; Bhavnani and Lee, 2021; California v. Cisco Systems, Inc., 2022) and remain challenging to misrepresent because of entrenched social networks (Beteille, 1969; Jodkha, 2017). Similar mechanisms have been documented in the US and U.K., where class or cultural signals affect perceptions of candidate suitability and belonging (Laurison and Friedman, 2019; Rivera, 2015).

5.6 Heterogeneity in Caste Penalties, Selectivity, and Promotion Outcomes: Empirical Facts & Stylized Model

In Sections 5.3–5.5, we demonstrated that most of the caste penalty in both earnings and success rates emerges at the personal interview stage, when caste becomes salient. Because we lack direct measures of candidate performance at each stage, any unobserved ability differences may remain concealed until later stages (Section 4.10).

This raises two natural questions. First, might these penalties appear earlier in *highly selective* jobs that recruit from the extreme right tail of the ability distribution? Second, do personal interviews reveal unobserved ability that could justify hiring differences across caste rather than merely reflecting discrimination?

By leveraging variation in job selectivity, we show that the empirical evidence does not sup-

port a "pure skill-revelation" explanation. The key insight is that variation in job selectivity provides a powerful test. If disadvantaged castes truly had lower performance scores at pre-interview screening rounds, such deficits would be most evident in highly selective jobs that recruit predominantly from the right tail. Yet, we observe negligible pre-interview caste gaps even in these jobs.

Moreover, once hired, disadvantaged castes exhibit higher promotion rates. One might argue this simply reflects inframarginal productivity differences—that is, more right-tail performers among the hired disadvantaged. However, examining promotion patterns across jobs of increasing selectivity helps distinguish between discrimination and ability differences: while jobs with identical hiring thresholds across castes would generate shrinking promotion advantages as they become more selective, we find these advantages remain persistent and stable across selectivity tiers.

Consistent with our focus on job selectivity, I present two empirical facts that exploit differences in hiring standards across jobs, and then develop a formal model demonstrating why these patterns indicate a stricter hiring threshold for disadvantaged castes. Fact 1 uses job-specific selectivity measures at each screening round and shows no caste gap in the success rate at pre-interview screening rounds across jobs of varying selectivity. Similarly, Fact 2 groups jobs into selectivity bins and shows stable and large promotion rate differences in favor of the hired disadvantaged across the selectivity distribution.

Fact 1 (**Conditional Success Rates Reveal No Early Penalty, but a Large Final Interview Gap**). Figure 5 plots, for each job, the estimated caste gap in *conditional* success rates (disadvantaged minus advantaged) on the y-axis against the cumulative proportion of applicants that job eliminates by a given screening round (defined as "cumulative cuts" in Section 4.6) on the x-axis. A negative gap indicates lower success rates for disadvantaged castes.

The y-axis values are calculated by estimating, for each job $j \in \mathcal{J}$, the following linear probability model:

$$\Pr(\text{Success}_{ijt}^{\text{Job Search Stage}}) = \alpha_j + \beta_j \times \text{Disadv. Caste}_i + \text{Controls}_i + \gamma_t + \epsilon_{ijt}, \tag{3}$$

where $i \in \mathcal{I}$ indexes students and γ_t denotes year fixed effects. The controls are characteristics such as college GPA, college entrance exam score, previous labor market experience, major, and degree.

Before personal interviews—across application reading, written tests, and group debates—the figure shows these caste gaps remain near zero, *even* among highly selective jobs that

eliminate a large share of applicants. Only at the final interview round does a distinct penalty for disadvantaged castes emerge, underscoring that late-stage "fit" assessments—rather than unobserved performance differences prior to interviews—drive the observed caste disparities.

As a robustness check, Online Appendix Figure OA.3 presents the *uncontrolled* success rates, which reveal that while disadvantaged castes face an initial penalty at the application reading stage, the penalty nearly disappears prior to personal interviews and then reemerges afterward—precisely where the controlled regressions show the gap intensifying in Fact 1.

Fact 2 (**Disadvantaged Hires Earn Comparably Higher Promotion Rates, Even in Highly Selective Jobs**)**.** The above patterns do not, however, rule out whether interviews are simply revealing other productivity-related traits in which disadvantaged castes genuinely lag. To distinguish unobserved ability differences from discrimination, we next turn to *post-hire* performance.

Using data from an employment-focused online platform, I track career trajectories for over 98% of the job fair sample and use promotion rates as a measure of job performance.[13] If disadvantaged castes truly had lower ability that was revealed at interviews, the relative performance of the hired disadvantaged should decline in more selective jobs that impose higher—but equal—standards across castes. Instead, as we show below, disadvantaged hires maintain their promotion advantage even in the most selective positions. This pattern aligns with outcome-based tests for discrimination developed in healthcare and algorithmic fairness research (Obermeyer et al., 2019; Chandra and Staiger, 2010, 2020; Abaluck et al., 2016), where higher success rates among marginalized groups, conditional on selection, are suggestive of discriminatory screening thresholds.

To compare caste differences in job performance, I estimate the following linear regression across job selectivity tiers:

$$\text{Outcome}_{ijt} = \alpha + \beta \times \text{Disadv. Caste}_i + \text{Controls}_i + k_j + \gamma_t + \epsilon_{ijt}, \tag{4}$$

where $i \in \mathcal{I}$ indexes students, $j \in \mathcal{J}$ indexes jobs, γ_t denotes year fixed effects, and k_j denotes job fixed effects. The controls are characteristics such as college GPA, college entrance exam score, previous labor market experience, major, and degree.

Table 7 shows the coefficients from these regressions:

- **Columns (1)–(2)** show that, pooling all jobs, disadvantaged hires have about a 20%

[13] The unlinked sample comprises just 1.5% of graduates. Section 4.9 details how these links are made.

higher probability of promotion relative to advantaged hires, despite having been penalized at the hiring stage.

- **Columns (3)–(8)** group jobs by their *selectivity* (final rejection rates) and reveal that a similar promotion advantage of about 20-25% for disadvantaged hires persists across tiers of increasing stringency—here, the top 75%, 50%, 40%, 30%, 20%, and 15% most selective jobs (e.g., jobs in the top 20% of selectivity reject at least 80% of initial applicants). In other words, even among candidates who survived highly stringent hiring standards, disadvantaged hires do *not* underperform on promotions.
- **Columns (9)–(10)** demonstrate similar patterns in both client-facing and non–client-facing roles, suggesting that the promotion advantage is not driven by a particular job function.

Promotion rate differences are likely conservative. Although these MNCs use globally standardized promotion criteria (Bloom et al., 2012), local managers exercise some discretion through subjective assessments of "leadership," "communication," and "potential" (Benson et al., 2019). Any resulting bias would tend to depress promotion rates for disadvantaged caste hires—making their strong observed promotion outcomes even more striking and reinforcing that they likely cleared a stricter hiring bar. The fact that disadvantaged caste hires achieve higher promotion rates despite possible managerial bias implies promotion rate differences are likely conservative.[14]

Complementing this argument is Online Appendix Table OA.9 showing that those who leave their first jobs at some point are typically more negatively selected with respect to GPAs if they came from advantaged castes, while promotions are weakly positively correlated with GPA (not shown).

Key Takeaway. Facts 1 and 2 challenge the notion that interviews merely catch up to latent skill deficits among disadvantaged castes. A large pre-interview performance deficit should have implied lower success rates for disadvantaged castes before interviews, especially among highly selective jobs. And, if interviews were accurately filtering out less-qualified candidates, we would expect similar or even lower promotion rates among those hired. Instead, disadvantaged hires consistently attain *higher* promotion rates—even in the most selective jobs. The promotion advantage does not shrink or equalize as jobs become more selective, a

[14] In reality, there is evidence of unfair treatment at work. A 2018 survey of South Asians in the US found that 67% of disadvantaged caste employees reported being treated unfairly at their American workplaces (California v. Cisco Systems, Inc., 2022).

discrepancy that is difficult to reconcile with explanations based solely on unobserved ability differences.

In what follows, I present a simple theoretical framework to rationalize these two empirical facts, showing how this evidence is consistent with discrimination in hiring—with personal interviews serving as the main mechanism that translates social status into hiring barriers.

5.6.1 Explaining the Empirical Facts: A Simple Model of Late-Stage Discrimination

This section presents a unified model showing how (i) the absence of any early shortfall for disadvantaged castes implies no genuine ability deficit before interviews, and (ii) multiple hiring thresholds across different jobs help distinguish unobserved ability differences revealed at the interview from discrimination.

We now present a single theoretical model that unifies both these elements.

Consider a set of jobs $j \in \mathcal{J}$. Each job has up to four sequential rounds $s \in \{1, 2, 3, 4\}$:

- Rounds $s = 1, 2, 3$ are pre-interview screening rounds (e.g. application, written tests, group debates)
- Round $s = 4$ is the personal interview

Groups and Screening. There are two castes (groups) $g \in \{0, 1\}$, where $g = 1$ denotes disadvantaged castes and $g = 0$ denotes advantaged castes. At each round s, employers observe a score $M_{s,i,g}$ for candidate i from group g.[15] Passing round s requires $M_{s,i,g} \geq \tau_{s,j}$ and also requires having passed all prior rounds.

For simplicity, we assume employers *perfectly observe* these screening scores. This assumption is not critical for the key insights derived from this model—variation in screening round selectivity and final hiring thresholds across jobs can help distinguish discrimination from ability differences even with imperfect information (see Online Appendix Section A.3 for this extension).

At the final interview ($s = 4$), the threshold $\tau_{4,j}^{(g)}$ may depend on group g since caste becomes

[15] We assume the performance distributions for each group are continuous with smooth densities, finite moments, and satisfy standard regularity conditions. These technical assumptions ensure well-behaved pass rates and expectations when calculating promotion probabilities. While these properties are implicit in our figures and analysis, they are not crucial for understanding the main results.

salient for nearly 85% of candidates at interviews (Section 5.5.2). Hence, any final-stage gap can reflect discrimination if $\tau_{4,j}^{(1)} > \tau_{4,j}^{(0)}$.

Monotonicity of thresholds. We assume

$$\tau_{4,j} \;\geq\; \tau_{3,j} \;\geq\; \tau_{2,j} \;\geq\; \tau_{1,j},$$

so each round is at least as demanding as the previous one.

Identification from exogenous variation in hiring thresholds across jobs. For our identification strategy, we assume that the underlying ability distributions are common across jobs. This assumption is crucial because it implies that differences in job thresholds τ_j reflect genuine variations in job selectivity rather than systematic differences in candidate quality.[16]

An intuitive way to think of this assumption is to imagine one large tub of ice cream representing the entire talent pool. Each job has its own "scoop"—its hiring threshold—to serve from that same tub. A more selective job uses a smaller scoop (higher threshold), thus hiring fewer but on average higher-scoring candidates. A less selective job uses a larger scoop (lower threshold), so it hires more candidates but at a lower average performance level. Crucially, everyone is scooping from essentially the same tub, which ensures that differences in scoop size truly reflect how "selective" a job is—rather than differences in the candidate pool itself.

Several institutional features support this assumption: (i) negligible application costs and broad applicant pools preclude sorting into jobs by caste (Section 4.7); (ii) admission quotas ensure balanced caste representation within majors and guarantee common support across GPA quantiles (Section 4.3 and Online Appendix Figure OA.1); (iii) caste is typically revealed only at interviews, thereby preventing selective screening in early rounds; and (iv) standardized recruitment protocols limit employers' ability to strategically manipulate cutoffs (Section 4).

Note that while *within* any given job the hiring cutoff is naturally correlated with candidate quality, our identification strategy exploits variation *across* jobs—where candidate quality is assumed comparable—so that differences in thresholds are driven solely by job selectivity. In our setting, each job (identified by its unique job ID) has its own hiring threshold, and

[16] Even if group-mean ability distributions vary slightly by job, as long as these deviations are mean zero or not systematically correlated with caste, the main results continue to hold.

because we assume that candidate ability distributions are comparable across jobs, the differences in these thresholds are exogenous and can be used as an independent instrument.

(A) Ruling Out Systematic Performance Differences Before Interviews

We rule out systematic caste differences in screening round performance before personal interviews.

Proposition 1 (No Early Gap Implies No Right-Tail Deficit Before Interviews)**.** Consider a job j with up to four rounds of screening, where the first three ($s = 1, 2, 3$) are pre-interview. If group 1 truly performed systematically worse, then for a sufficiently large threshold $\tau_{s,j}$ at some pre-interview stage $s < 4$, the pass-rate gap

$$\Delta_{s,j} \;=\; P(\text{Pass}_{s,i,j}|g = 1) \;-\; P(\text{Pass}_{s,i,j}|g = 0) \;<\; 0$$

must arise. Therefore, if $\Delta_{s,j} \geq 0$ even under high thresholds—essentially focusing on the top tail of each group's performance distribution—this rules out a meaningful right-tail performance deficit for group 1 in earlier rounds.

Sketch of Proof. We illustrate how variation in screening thresholds helps identify any systematic performance differences before interviews.

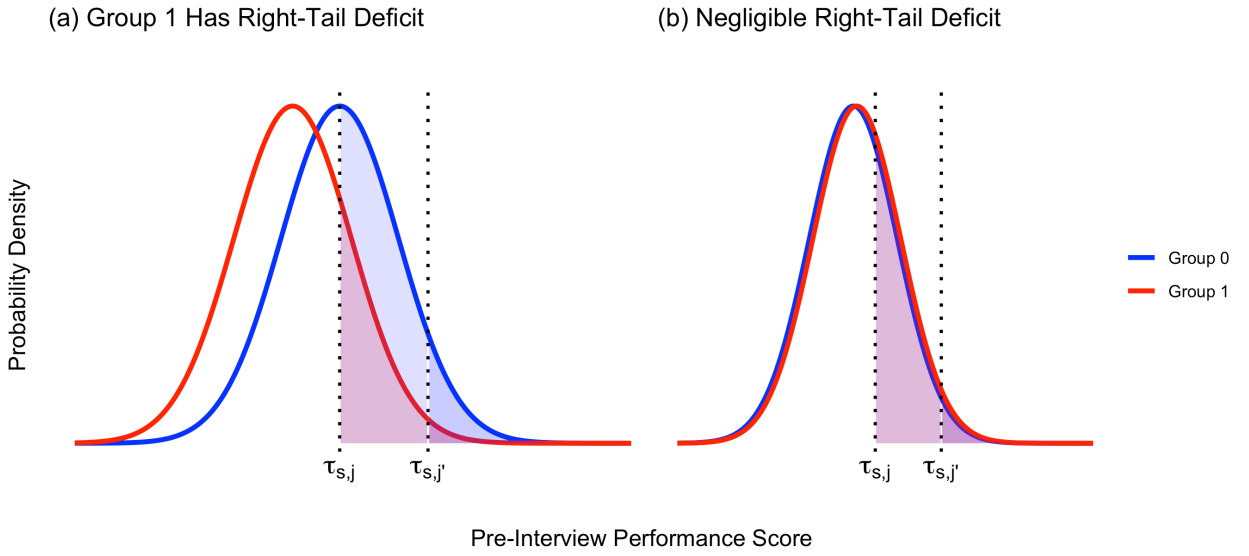
Figure 3 contrasts two scenarios. In each panel, the blue (solid) and red curves depict two groups' performance distributions in pre-interview screening rounds. Recall that $g = 1$ denotes the disadvantaged group and $g = 0$ the advantaged group. We shade the portion of each group's distribution above the vertical dashed line, representing the screening cutoff at that job. The contrast between panels reveals why selective jobs provide a powerful test for pre-interview performance differences:

- **Panel (a): Hypothetical Right-Tail Deficit.** Assume that for each group g, the screening score distribution $\varphi_g(m)$ has a regularly decaying tail; for example, for sufficiently large m, $P(M_{s,i,g} \geq m) \sim C_g\, m^{-\alpha}$ for some constants $C_g > 0$ and $\alpha > 1$. This condition implies that as the screening threshold increases (i.e., in more selective jobs), the conditional pass rate declines at a rate determined by the tail density. Consequently, if Group 1 has a thinner right tail than Group 0 (i.e., a smaller C_1 or larger effective decay rate), the red "passing region" would shrink *faster* than the blue, leading to increasingly *negative* pass-rate gaps in more selective jobs.
- **Panel (b): No Meaningful Deficit.** Fact 1 instead shows no systematic shortfall for Group 1 across the pre-interview screening performance distribution. Even at high

cutoffs in selective jobs, both groups maintain similar pass rates. This pattern holds across application reading, written tests, and group debates.

The observation of similar pass rates across groups in highly selective jobs (Panel b) thus rules out systematic pre-interview performance differences.

Figure 3: **Using Variation in Screening Thresholds to Test for Pre-Interview Performance Differences**



Notes: Figure 3 illustrates how selective screening would reveal systematic pre-interview performance differences. Panel (a) shows how a hypothetical right-tail deficit for Group 1 would become more apparent at higher cutoffs. Panel (b) shows the likely distribution of pre-interview performance scores implied by Fact 1, with no meaningful deficit. $\tau_{s,j}$ denotes screening threshold at stage $s < 4$ for job j (j' represents a more selective job).

□

Key Takeaway. Fact 1 shows that early-stage disadvantages remain negligible before interviews, thereby corroborating Proposition 1. Consequently, any late-stage gap at $s = 4$ must stem either from newly revealed ability signals at the interview or from discrimination.

(B) Using Variation in Hiring Thresholds to Test for Discrimination

We now use variation in hiring thresholds across jobs to show that the persistently higher promotion rates for disadvantaged castes across selectivity tiers are most consistent with hiring discrimination, rather than unobserved ability differences.

We begin by linking screening performance to job performance. Maintaining our perfect information assumption, we define job performance as the average screening score of accepted candidates. Let $\bar{M}_{i,g} = \frac{1}{4}\sum_{s=1}^{4} M_{s,i,g}$ denote the average screening score for candidate i from

group g (any monotonic link $f(.)$ between screening round performance and job performance works, but we choose to keep things simple). A candidate is promoted if $\bar{M}_{i,g}$ exceeds a job-specific threshold κ_j.

For jobs in the top $(1-q)$ fraction by selectivity (for instance, if $q = 0.8$, then we are conditioning on the top 20% of jobs by selectivity), we distinguish between two types of thresholds: the *selectivity-tier threshold* $\tau^q_{4,j}$, which captures the job-level standard for jobs in the top $(1-q)$ fraction, and the *group-specific threshold* $\tau^{(g)}_{4,j}$, which may vary across groups within each job. Under discrimination, $\tau^{(1)}_{4,j} > \tau^{(0)}_{4,j}$; if no discrimination is present, these thresholds coincide (i.e., $\tau^{(1)}_{4,j} = \tau^{(0)}_{4,j} = \tau^q_{4,j}$).

Define the average performance among hired candidates for group g as

$$\bar{Y}^{(g)}_j = E\left[\bar{M}_{i,g} \mid M_{4,i,g} \geq \tau^{(g)}_{4,j}\right],$$

where the expectation is taken over the ex-ante distribution of screening scores. Then, the promotion gap at tier q is given by

$$\Delta^{\text{Prom}}(q) = \bar{Y}^{(1)}_j - \bar{Y}^{(0)}_j.$$

The key insight is that as the hiring threshold increases, the *average* screening score of accepted candidates converges to the quality of the *marginal* candidate at the threshold. Thus, if both groups face the same threshold, the difference in average screening scores—and hence the promotion gap—would shrink toward zero. Conversely, if disadvantaged candidates face a higher threshold (i.e., $\tau^{(1)}_{4,j} = \tau^{(0)}_{4,j} + \delta$ for some $\delta > 0$), then the marginal candidate for Group 1 is of higher quality than that for Group 0, and the average performance gap remains roughly constant.

Proposition 2 (Stable Promotion Gaps Across Selectivity Tiers Indicate Discrimination)**.** Let J_q denote the set of jobs in the top $(1-q)$ fraction by selectivity with associated thresholds $\tau^q_{4,j}$. Suppose that for jobs in J_q, the average screening score among accepted candidates for group g is given by

$$\bar{Y}^{(g)}_j = E\left[\bar{M}_{i,g} \mid M_{4,i,g} \geq \tau^{(g)}_{4,j}\right].$$

If both groups face the same hiring threshold (i.e., $\tau^{(1)}_{4,j} = \tau^{(0)}_{4,j}$), then as $q \uparrow$ (i.e., as jobs become more selective and $\tau^q_{4,j}$ increases), the average screening scores of accepted candidates in both groups converge to the common threshold, implying that $\Delta^{\text{Prom}}(q) \to 0$.

Conversely, if discrimination is present so that $\tau^{(1)}_{4,j} = \tau^{(0)}_{4,j} + \delta$ for some $\delta > 0$, then the average

screening score for Group 1 converges to a value that is δ higher than that for Group 0, and the promotion gap remains approximately constant at δ across selectivity tiers.

Sketch of Proof. Our aim is to illustrate how variation in hiring thresholds distinguishes discrimination from ability differences.

Figure 4 contrasts two scenarios, with blue and red curves depicting the average screening round performance scores for the two groups. The red curve (Group 1) is shifted rightward relative to the blue curve (Group 0), consistent with both our assumption that average screening round performance is positively correlated with job performance and our empirical finding that disadvantaged castes have higher promotion rates post-hire. The vertical dashed lines represent job-level cutoffs, and the shaded regions above these thresholds denote the hired populations.

- **Panel (a): No Discrimination.** In the absence of discrimination, both groups face a common cutoff, say $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} = \tau$. Under the regularity condition that each group's screening score distribution decays regularly (i.e., for sufficiently large m, $P(M_{s,i,g} \geq m) \sim C_g\, m^{-\alpha}$ for some $C_g > 0$ and $\alpha > 1$), it follows that

$$\lim_{\tau \to \infty} \left(E[\bar{M}_{i,g} \mid M_{4,i,g} \geq \tau] - \tau \right) = 0.$$

  Hence, as the common cutoff shifts rightward in more selective jobs, the average screening scores of accepted candidates in both groups converge to the cutoff value (i.e., the performance level of the marginal candidate), and the performance gap diminishes toward zero.

- **Panel (b): Discrimination.** If discrimination is present, then $\tau_{4,j}^{(1)} > \tau_{4,j}^{(0)}$ (note that when no discrimination is present, these thresholds coincide with the selectivity-tier threshold, i.e., $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} = \tau_{4,j}^{q}$).

  Suppose that Group 1 faces a higher cutoff than Group 0 by a constant $\delta > 0$ (i.e., $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} + \delta$). Under the regularity condition that for each group g,

$$\lim_{\tau \to \infty} \Big(E[\bar{M}_{i,g} \mid M_{4,i,g} \geq \tau] - \tau \Big) = 0,$$

  it follows that as the threshold increases,

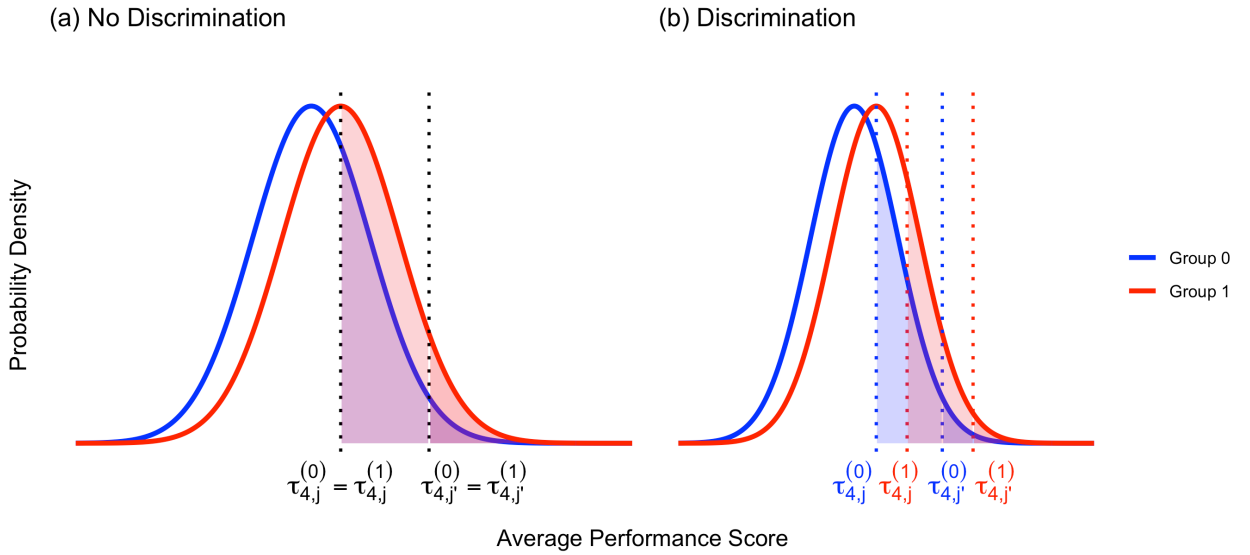
$$\lim_{\tau \to \infty} \Big(E[\bar{M}_{i,1} \mid M_{4,i,1} \geq \tau_{4,j}^{(1)}] - E[\bar{M}_{i,0} \mid M_{4,i,0} \geq \tau_{4,j}^{(0)}] \Big) = \delta.$$

  In other words, the average screening score for Group 1 converges to a value that is δ

higher than that for Group 0, so the performance gap remains approximately constant at δ even as the threshold increases.

Since a candidate is promoted if $\bar{M}_{i,g}$ exceeds a job-specific threshold κ_j, the asymptotic behavior of the screening score distributions implies that if both groups face the same threshold, the promotion gap $\Delta^{\text{Prom}}(q)$ shrinks toward zero, whereas a constant gap in thresholds (indicative of discrimination) preserves a stable promotion gap.

Figure 4: **Using Variation in Hiring Thresholds To Test for Discrimination**



Notes: Figure 4 depicts two scenarios. In panel (a), a common cutoff for both groups shrinks their hired-performance gap at higher cutoffs. In panel (b), a constant δ gap in cutoffs preserves a stable performance advantage for Group 1. Here, $\tau^{(g)}_{4,j}$ denotes group g's interview threshold at job j, with j' representing a more selective job.

□

No requirement that all hires are marginal. Importantly, this argument does not require every hire to be exactly at the threshold. Even if many hires are well above the threshold, examining patterns across selectivity tiers identifies systematic differences in hiring standards—the fundamental distinction between discrimination and genuine ability differences.

Key Takeaway. Fact 2 offers compelling evidence for the discrimination scenario outlined in Proposition 2. The stable promotion gap across selectivity tiers—where disadvantaged hires maintain a roughly 20% advantage—is inconsistent with an ability-differences explanation, which would predict a diminishing gap at higher thresholds.

Summary. The absence of caste penalties in even the most selective jobs' early rounds rules out meaningful ability deficits before interviews (see Fact 1 and Proposition 1). More-

over, when disparities do appear at interviews, their consistent pattern across jobs distinguishes their origin: the stable 20% promotion advantage observed for disadvantaged castes across all selectivity tiers (Table 7) is precisely what we would expect from discriminatory hiring thresholds (see Fact 2 and Proposition 2), rather than from unobserved ability differences.

5.7 Preferred Interpretation of the Caste Penalty

Given the evidence that employers directly penalize caste, my preferred interpretation is that hiring disparities stem from *disparate treatment*—i.e., discrimination based on caste. Caste-based discrimination is explicitly illegal under Article 15 of the Indian Constitution (Ministry of Law and Justice, Government of India, 2021). Disparate impact could also be at play (for instance, if "fit" criteria implicitly favor certain extracurricular backgrounds), and the caste penalty embeds all such alternative rationalizations.

5.8 What Happens During Personal Interviews?

The paper's main goal is not to pin down the exact content of the conversations at personal interviews. It is reasonable to assume that interviews screen on multiple attributes—such as background, extracurriculars, and "fit"—often beyond what is visible on paper. Given the findings in Sections 5.3–5.6, however, *any* additional diagnostic information gleaned at interviews cannot justify the observed hiring disparities.

This section provides suggestive evidence on personal-interview practices for elite multinationals. There are three key takeaways from this section. First, personal interviews are inherently subjective. Second, such interviews serve as avenues for assessments of candidate "fit," which includes both legitimate skill- or role-related evaluations and more intangible cultural or background alignment that hiring managers consider relevant during the interview process. Third, cultural similarities between interviewers and candidates significantly influence hiring decisions at elite firms, potentially perpetuating socioeconomic homogeneity. This section should *not* be interpreted as providing dispositive evidence on the content of personal interviews in my setting. I do not have access to interview transcripts. Nevertheless, it is important to emphasize that evidence from other fields examining hiring practices in high-stakes settings has shown that personal interviews are often subjective and serve as common avenues for assessments of many "soft" characteristics, including candidate background and fit, that are strongly correlated with social class and privilege.

Personal interviews at the US operations of elite firms. As Lamont (2009) observes, cultural qualities tend to be more salient in settings where differences in technical quality are minimized, such as elite hiring contexts. Similarly, Lamont (1992) has established that cultural signals play a particularly important role among elites. Within this theoretical framework, comprehensive studies by Rivera (2012, 2015) of hiring professionals at elite multinationals provide one of the first and only systematic examinations of the content of personal interviews. Through 120 interviews with professionals involved in hiring at top-tier law firms, investment banks, and management consulting firms, Rivera found that "cultural fit" was not merely an informal preference but a formal evaluative criterion mandated by these organizations and embraced by individual evaluators.

Significantly, more than half of Rivera's respondents reported that fit was the most important criterion they used to assess candidates in job interviews, rating it above analytical thinking and communication skills. This emphasis varied by industry, with over 70% of law firm partners, 65% of investment bankers, and 40% of consultants ranking fit as their top evaluative criterion (Rivera, 2012). This focus on fit remained strong even in consulting firms, which employed the most technical interview formats through case-based business questions.

In elite firms surveyed by Rivera (2012), roughly 80% of evaluators used a heuristic known as the "airplane test" to assess candidate fit. As one investment banking director explained:

> One of my main criteria is what I call the stranded in the airport test. Would I want to be stuck in an airport in Minneapolis in a snowstorm with them? And if I'm on a business trip for two days and I have to have dinner with them, is it the kind of person I enjoy hanging with? And you also have to have some basic criteria, skills and smarts or whatever, but you know, but if they meet that test, it's most important for me.

Interviewers actively sought biographical commonalities with candidates, as illustrated by one attorney:

> I usually start an interview by saying, "Tell me about yourself." When I get asked that, I talk about where I'm from, where I was raised, and then my background. A not-good way to start is with law school. I want to hear your life story. Hopefully there's something more interesting about your life than deciding to go to law school... When they tell me about their background, it's easier to find things in common... Maybe... they're from Seattle and I've been to Seattle. We

can talk about that and develop a connection.

What's particularly noteworthy is that these evaluations of fit were not merely about work styles or professional competencies but specifically focused on candidates' extraprofessional and leisure activities—dimensions often shaped by socioeconomic background and access to resources. Rivera distinguishes fit from the communication skills required in client-facing professions, which evaluators grouped into the separate category of "polish" or "presence." As one consultant explained:

> When you are judging someone [to see] if you want to put him in front of a client, the question is do they conduct themselves professionally... You need someone who speaks in a way that earns your trust, who presents their opinion respectfully but also convincingly... But in terms of "fit," it's someone that we want on our case team... You want someone that makes you feel comfortable, that you enjoy hanging out with, can maintain a cool head when times are tough and make tough times kind of fun.

Once candidates demonstrated a baseline of competence, perceptions of fit rather than absolute technical performance routinely drove assessments. One consulting manager recalled:

> On the fit side, I wrote [on the evaluation form]... 'Will quickly become everyone's best friend.'... That's what I call a good fit. But quite frankly, his case performance wasn't the best. But because his personality and presence were so strong, I forwarded him on [to second-round interviews].

The emphasis on cultural similarities was also reflected in how evaluators justified their hiring decisions. Rivera found that evaluators constructed and assessed merit in their own image, believing that candidates who were culturally similar to them would have a higher likelihood of job success (Rivera, 2015). Banking recruiting head Stephanie summarized this tendency bluntly: "You are basically hiring yourself. This is not an objective process" (Rivera, 2012).

Furthermore, the specific types of cultural similarities valued in elite firms were not merely about personality traits but about specific leisure pursuits and extracurricular interests that signal class position. Rivera found that evaluators tended to favor extracurricular activities associated with the white upper-middle class and that were acquired through intense, prolonged investment of material and temporal resources (Rivera, 2015). Employers used these "fine distinctions" to screen résumés and compile interview pools, with activities like varsity sports, playing the cello, or enjoying particular cultural pursuits serving as critical signals of

fit and, by extension, class membership.

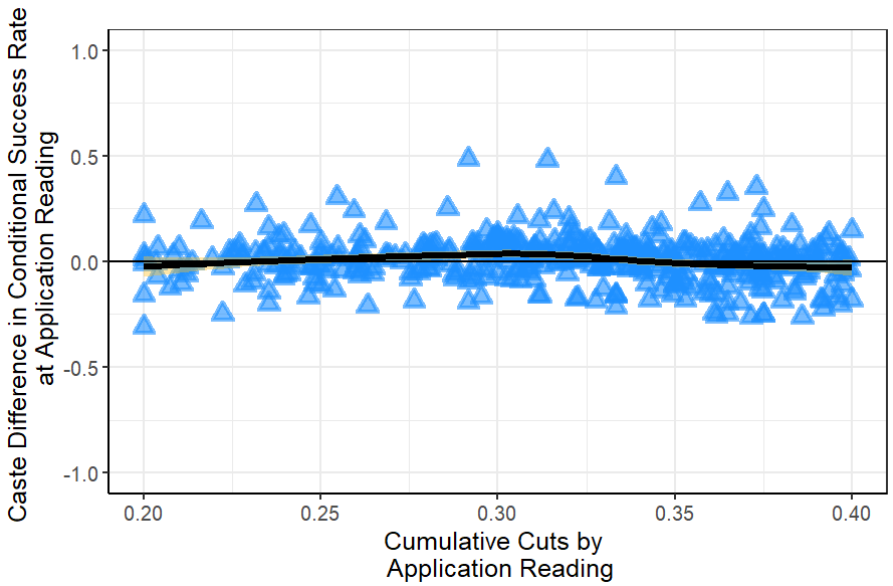
While Rivera's research focuses on elite multinational firms in the United States, similar patterns of subjective evaluation and class-based filtering emerge in other high-stakes selection processes worldwide.

Personal interviews in India's elite civil service. The emphasis on subjective assessments is also mirrored in India's elite bureaucracy. This dynamic appears to directly impact candidates from marginalized backgrounds in the Indian context. A former UPSC chairman candidly acknowledged the biases introduced through the interview process: "A candidate from a rural background and educated at a small place finds it difficult to compete in communication skills before the interview panel with those who are from cities, and have been educated in a better atmosphere" (Bhavnani and Lee, 2021, p. 27). The result is that even high-performing candidates from disadvantaged groups face significant barriers in the subjective evaluation phase, effectively creating a "cultural tax" on advancement.

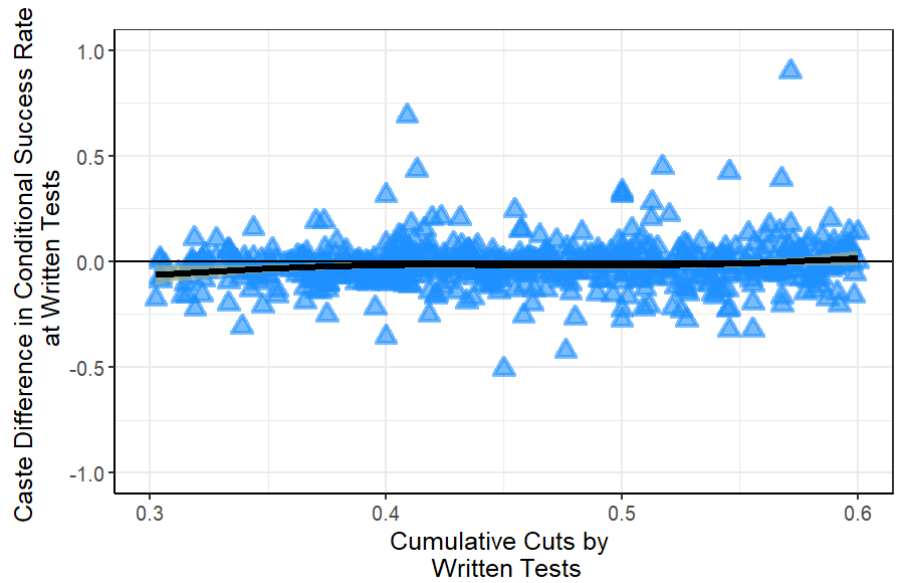
These IAS findings conceptually replicate the patterns seen in multinational firms, reinforcing a broader theme: subjective, "fit"-focused screening can systematically harm lower-status candidates without reliably predicting job success. In both the multinational and IAS contexts, final-stage interviews penalize social class. Meanwhile, promotion rates or other objective performance indicators suggest that many highly capable, lower-class candidates get sidelined for reasons unrelated to actual productivity or skill. This recurring pattern across both private corporate hiring and government bureaucratic recruitment suggests that subjective, background-sensitive interviews risk undermining the very organizational goals of securing the best talent, regardless of social status.

These patterns reveal how personal interviews at elite firms serve as mechanisms that can systematically advantage candidates who share cultural markers with their evaluators, potentially creating barriers for candidates from different socioeconomic backgrounds regardless of their technical qualifications or potential for success. The findings suggest that what appears on the surface as "meritocratic selection" may in fact function as a form of class-based filtering that reproduces existing social hierarchies under the guise of objective assessment.

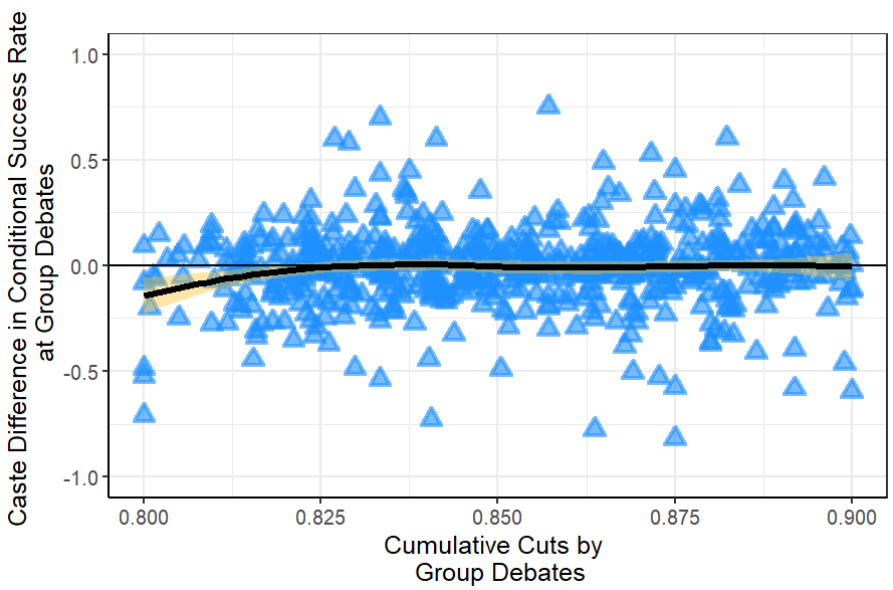
Figure 5: **Caste Penalty in the Conditional Success Rate At Each Screening Round**



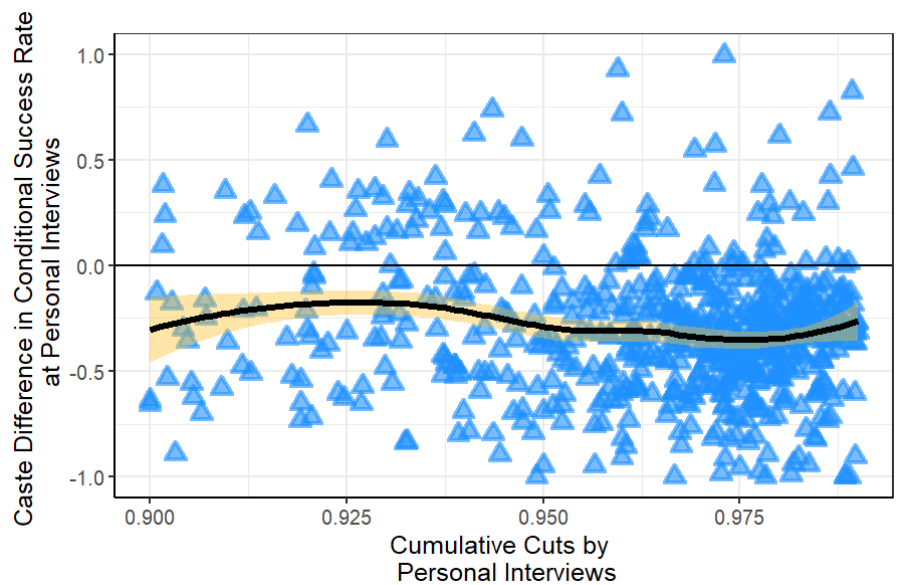
(a) Caste Penalty in Conditional Success Rates versus Cumulative Cuts by Application Reading



(b) Caste Penalty in Conditional Success Rates versus Cumulative Cuts by Written Tests



(c) Caste Penalty in Conditional Success Rates versus Cumulative Cuts by Group Debates



(d) Caste Penalty in Conditional Success Rates versus Cumulative Cuts by Personal Interviews

Notes: Figure 5 shows the caste difference in the conditional (i.e., with controls) success rate against the cumulative cuts made by jobs by the end of each screening round. Each blue triangle denotes a job. Smoothed conditional means (thick black lines) are calculated using the "loess" function in R. The shaded areas in yellow denote 95% confidence intervals. The thin solid horizontal line denotes a y-intercept of 0. The x-axis shows the cumulative proportion of students cut by jobs by the end of each screening round. The y-values are calculated by estimating, *for each job*, a standard linear probability model with controls and then plotting the coefficient on disadvantaged caste from these regressions on the y-axis. Each regression includes controls for student characteristics such as degree, major, college GPA, entrance exam score, and cohort year.

Table 7: **Caste Differences in Promotion Rate**

	Dependent variable:									
	Promoted	Promoted	Promoted at Top 75%	Promoted at Top 50%	Promoted at Top 40%	Promoted at Top 30%	Promoted at Top 20%	Promoted at Top 15%	Promoted at Client Facing	Promoted at Non-Client Facing
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Disadv. Caste (OLS)	0.137***	0.096***	0.104***	0.112***	0.136***	0.147***	0.190**	0.204**	0.107**	0.082**
	(0.025)	(0.026)	(0.029)	(0.038)	(0.043)	(0.054)	(0.073)	(0.084)	(0.044)	(0.032)
Disadv. Caste (Oster)	0.142***	0.093***	0.108***	0.122***	0.141***	0.157**	0.196**	0.206**	0.106***	0.075**
	(0.026)	(0.028)	(0.031)	(0.037)	(0.041)	(0.058)	(0.082)	(0.082)	(0.039)	(0.036)
Tenure Duration		0.084***	0.082***	0.076***	0.074***	0.068***	0.054*	0.079**	0.065***	0.097***
		(0.012)	(0.014)	(0.019)	(0.021)	(0.025)	(0.030)	(0.037)	(0.021)	(0.015)
Mean Y	0.684	0.684	0.667	0.668	0.676	0.711	0.711	0.706	0.687	0.682
Observations	2,489	2,489	2,012	1,246	997	748	502	374	835	1,654
Adjusted R^2	0.183	0.205	0.203	0.216	0.211	0.182	0.192	0.189	0.194	0.214
Job FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Other Controls		✓	✓	✓	✓	✓	✓	✓	✓	✓

Notes: Table 7 shows caste differences in job performance measures among students who graduated with full-time jobs through the college's on-campus job fair. All of these students had to succeed in the final personal interview round to obtain jobs. Column (1) reports caste differences in promotion rates only with job and year fixed effects. Column (2) reports the same regression as Column (1) but with all other controls such as GPA, major, degree, and job tenure. Columns (3)-(8) show caste differences in promotion rates with controls among jobs at different quantiles of "job selectivity," proxied by the cumulative proportion of applicants cut by a given job by the end of the final screening round. Columns (9)-(10) show caste differences in promotion rates with controls among client-facing and non-client-facing jobs, respectively. Bias-adjusted coefficients are calculated using Oster (2019). Following Oster's recommended benchmarks, I set $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$, implying that unobservables are assumed to have the same degree of selection as observables, and the total explanatory power can increase by up to 30% beyond the observed R^2. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

6 Alternative Explanations

In this section, I consider possible alternative explanations for why caste disparities in hiring outcomes might predominantly emerge after personal interviews. While each explanation may offer insights, none provides a compelling account of the observed patterns.

1. **Early "fairness" by advancing weaker candidates.** One possibility is that employers lower their thresholds in the initial stages to appear fair. If so, disadvantaged castes, who have weaker academic credentials (e.g., GPA), should show higher conditional success rates relative to observably comparable advantaged castes before the interview. This effect would be especially pronounced at highly selective jobs. Yet, Figure 5 shows no such advantage (or penalty) at any pre-interview round. Although disadvantaged castes with clearly identifiable surnames face a penalty at the application stage (Figure 2), it contradicts the notion of early "fairness" for the sake of virtue signaling, as students employers would most want to give a bump to are those with obvious caste names. Further, India lacks a private-sector regulatory body akin to the EEOC, and many elite MNCs perceive caste-conscious hiring as "anti-merit" (Jodhka and Newman, 2007; Newman and Thorat, 2012). If anything, employers seeking to appear fair would do so at the final hiring stage, not before.

2. **"Delayed preference" for caste in interviews.** Another hypothesis is that employers delay penalizing caste until interviews. However, given the limited institutional pressure to improve caste diversity at earlier stages, there is no clear incentive to "save" discriminatory actions for later. If employers already know candidates' caste, there is little reason to postpone penalizing it. Indeed, Figure 2 shows that for students with identifiable surnames, the penalty appears immediately at the application stage. Moreover, if interviews were simply more subjective, we would expect group debates—another subjective round—to also reveal disparities, but the data do not support that. Group debates provide less scope to identify caste (Section 5.5.3), which aligns with the absence of a caste penalty at this stage.

3. **Inframarginal caste differences in performance before interviews.** This would imply a visible caste penalty before interviews, especially in highly selective jobs. Figure 5 shows no evidence of such a penalty, ruling out this hypothesis.

4. **Minimal student attrition at each stage.** Each job has four screening rounds, eliminating candidates after each (Figures 5, OA.3, and Table 4). There is no major student attrition unaccounted for in these data.

5. **Better outside options for advantaged castes.** Another explanation might be that advantaged castes withdraw early for better outside options. However, fewer

than 5% of students leave the main placement process by accepting summer internship offers (Section 3.2). Once registered, students cannot simultaneously search "offline" without losing placement office support for full-time job search. Hence, there is little scope for parallel external job searches during the process.

6. **Differences in socio-emotional skills.** If disadvantaged castes lacked socio-emotional abilities, we would see a notable penalty in group debates—another highly subjective round designed to assess communication and teamwork. Panels (c) of Figures 5 and OA.3 show that no such penalty arises, even in more selective jobs.

7. **Missing data because of employers or students skipping screening rounds.** The placement office ensures that employers conduct all four screening rounds and that all students attend the screening rounds they qualify for. It also collects comprehensive data at every stage (Section 4.2). Moreover, exit data confirms that job outcomes align with placement records (Section 5.2). Therefore, missing data cannot explain the observed hiring disparities.

8. **Negotiation advantages for advantaged castes.** The salaries and amenities offered during on-campus placements are non-negotiable (Section 3.2). The placement office verifies that final compensation matches the advertised package, and exit data confirms that job-getters commence with identical salaries months later (Section 5.2). Thus, negotiation disparities cannot explain the observed hiring disparities.

9. **Competition from government jobs.** Government positions with caste-based quotas comprise under 2% of all jobs in this sample (Section 4.5). Moreover, the list of participating firms is publicly available on the college portal, so private-sector employers do not implicitly compete with a large public-sector presence.

10. **Caste differences in job preferences.** Figure 1 shows that differences in job preferences across castes explain only a small portion of the overall initial hiring gap. These preferences thus cannot account for the dominant caste penalty that materializes at personal interviews.

11. **Preferential treatment for disadvantaged castes during promotions:** Few MNCs prioritize caste or class in diversity initiatives—with under 10% of top US companies even mentioning "class" (Stansbury and Rodriguez, 2024). Moreover, less than 3% of firms on India's premier stock exchange collect caste-related employment data (BusinessLine, 2018), indicating caste-conscious policies are rare, even among Indian managers. Evidence also shows workplace caste-bias typically works against disadvantaged castes, not in their favor (California v. Cisco Systems, Inc., 2022).

7 External Validity

A potential concern is that this study focuses on hiring from one elite Indian college for entry-level jobs at multinationals in India. Nonetheless, as I argue below, these multinationals follow comparable hiring practices across dozens of similar colleges and in global offices, so the findings illuminate broader processes that operate in high-stakes, elite recruitment across diverse contexts.

Caste as a well-defined lens for studying systemic inequality. In most global labor markets, examining discrimination based on social position poses significant measurement hurdles. Status markers such as income, accent, or social networks can overlap, complicating efforts to pinpoint precisely *how* and *when* discrimination arises. A first-generation college graduate with comparable academic credentials to a legacy student may face barriers rooted in institutional familiarity (à la Bourdieu, 1986).

Following Bourdieu's framework and broader sociological accounts (Lamont and Lareau, 1988; Rivera, 2015), social class comprises economic capital (material resources), social capital (networks and connections), and cultural capital (education, accent, mannerisms). In Western contexts, these forms of capital often manifest independently—a wealthy background (economic capital) does not automatically confer elite networks (social capital) or refined mannerisms (cultural capital). By contrast, in India, the presence of "poor advantaged-caste" or "wealthy disadvantaged-caste" families does not erode caste's grip on perceived status, as entrenched hierarchies typically overshadow such economic exceptions (Beteille, 1965, 1969).

India's caste system thus offers a uniquely comprehensive stratification marker that consolidates economic, cultural, and social capital into a single, observable measure. This clarity allows researchers to isolate the precise point at which informal "fit" screenings penalize candidates based on social status. In contrast to Western contexts—where race and gender are well-studied but "class" remains harder to pin down—caste in India merges these perceived status signals into a single, recognized hierarchy. By leveraging these measurable caste categories, this paper pinpoints the presence of discrimination and the specific mechanism through which it manifests: subjective "fit"-focused personal interviews.

Moreover, because caste provides a signal of social hierarchy, the results here are instruc-

tive for understanding how *any* consistently identifiable status marker—whether commonly perceived as visible (e.g., race) or more concealed (e.g., LGBTQ+ identity)—can be targeted through late-stage, subjective hiring steps.

Relevance across elite institutions. Nearly all top Indian colleges hold structured on-campus job fairs with analogous rules, screening protocols, and pools of elite employers (Sections 3.3 and 4.2). These features reinforce the generalizability of this study within India. These firms source their elite entry-level hires from a similar network of about 30-40 institutions each year, each with roughly 1500 students, suggesting these practices affect approximately 50,000 students annually.

Moving beyond India's borders, because these standardized, multi-round processes culminate in "fit"-focused interviews, they extend naturally to large-scale recruitment worldwide (Rivera, 2015, 2012). Many leading multinationals implement similar multi-round hiring practices across offices, so the processes observed here capture essential features of how global firms recruit talent in high-stakes environments.

Even in contexts where social status is gauged through less clearly defined markers than caste, final-stage subjective assessments often play a similar gatekeeping role for class or other underrepresented identities. Just as caste consolidates multiple status signals in India, these elite recruitment processes worldwide often blend evaluations of economic resources, social networks, and cultural attributes into subjective assessments of "fit." Accordingly, the core mechanism—discretionary interviews focusing on "fit"—applies to many high-skill labor markets beyond India, where recruitment decisions often rely on interpersonal judgments.

Global mechanisms and relevance. Although caste is specific to India, the underlying phenomenon of subjective criteria revealing and penalizing social standing is global in scope. Many employers worldwide depend on informal conversations about "cultural fit," inadvertently screening out underrepresented talent. Evidence from the UK and the US shows that subtle proxies for social class—such as accent or schooling—play similar gatekeeping roles, echoing the dynamics uncovered here (Laurison and Friedman, 2019; Rivera, 2015).

At the same time, parallels can be drawn to "test-optional" admissions or holistic evaluations in higher education (Sacerdote et al., 2025; Chetty et al., 2023), which similarly rely on subjective or "soft" information. These processes, though potentially well-intentioned,

may disadvantage high-talent individuals from lower-status backgrounds because they introduce intangible assessments—akin to late-stage "fit" conversations in hiring—that inadvertently reinforce elite advantages.

Moreover, as "traditional" status signals (e.g., pronounced accents or surnames) diminish in urban or highly educated settings (Clark, 2015; Wiltshire, 2020), status often becomes clear only through deeper, interpersonal interactions (Bourdieu, 1986). As these obvious cues recede, final-round conversations—such as those in personal interviews—grow *grow even more pivotal* for exposing concealed aspects of the advantages or disadvantages of social background in hiring. This evolution highlights the value of step-by-step data for pinpointing when and how discrimination arises: even where visible markers are diluted, subjective hiring conversations remain a potent mechanism for reproducing social inequality.

Erosion of affirmative action gains. This paper also highlights the fragility of affirmative action benefits during private-sector hiring. In India, government policies reserve seats in public colleges for disadvantaged castes, yet private firms—particularly global multinationals—are not bound by those provisions and often exclude social status from their DEI efforts, with less than 10% of firms in the Fortune 500 or the S&P 500 collecting information on social class or background (Stansbury and Rodriguez, 2024).

Although affirmative action helps disadvantaged students enter top institutions, the analysis shows how their progress can stall at the hiring stage through "fit" interviews that reveal caste. While recent research finds that partial or improper enforcement of affirmative action can disadvantage protected groups (Yenmez and Sönmez, 2022), this study pinpoints a *distinct* mechanism: even when policy might effectively admit disadvantaged students into elite colleges, subjective "fit" screens can reinstate social barriers downstream. Furthermore, these patterns speak to the broader tension noted by Chetty et al. (2023) and Sacerdote et al. (2025): when "soft" evaluative measures replace more standardized or transparent assessments, class advantages can be amplified rather than mitigated. These findings echo concerns globally that the labor-market gains of targeted educational policies risk erosion once informal screening and social biases come into play (Rivera, 2015).

In sum, although this study focuses on one prominent Indian college and multinational firms, it provides the *first systematic evidence* quantifying both the salience and the

source of social status (class) discrimination in elite hiring. By collecting step-by-step recruitment data from each screening round, we show that nearly 90% of hiring disparities materialize in subjective final interviews, illustrating how discretionary "fit" assessments penalize social status, thereby undermining both diversity and talent discovery. Whether in India or elsewhere, these informal evaluations often marginalize highly qualified candidates whose social status is perceived as incongruent with the firm's "culture," thereby perpetuating inequality.

In an era of global competition for high-skilled workers, recognizing and reforming these "final hurdle" practices is crucial for both achieving equity and optimizing talent allocation. Structured, transparent screening methods could help curb the inefficiencies created when gatekeepers rely on open-ended "fit" judgments. Understanding these processes is essential for future research and policy interventions aimed at fostering more meritocratic hiring standards worldwide.

8 Conclusion

This paper provides the first large-scale, quantitative account of *where* status-based or class discrimination emerges in elite hiring and *how much* it matters. To do so, it leverages a unique setting that measures social status (caste) and tracks applicants through four screening rounds at major multinational firms. Late-stage personal interviews assessing "fit" are the primary mechanism for revealing and penalizing social status, accounting for nearly 90% of the resulting disparities.

By focusing on subtle social cues—from extracurriculars to conversational style—recruiters effectively reintroduce status-based judgments at the moment of hire. Even highly credentialed candidates who have overcome significant educational barriers can suddenly be deemed unsuitable if they fail an amorphous "culture check." This dynamic extends across contexts: whether in India, Europe, or the United States, subjective evaluative steps can impose similar penalties based on race, sexual orientation, or class background, revealing a universal vulnerability in how late-stage assessments are conducted (Rivera, 2015; Laurison and Friedman, 2019; Stansbury and Rodriguez, 2024).

These findings also resonate with the dual-process cognition framework (Kahneman, 2011), which suggests that quick, intuitive judgments in high-stakes interview settings

can override more reflective evaluations of true ability. Similarly, Sarsons et al. (2021) and related work illustrate how subjective assessments can distort recognition of actual performance, reinforcing biases that more structured methods might mitigate.

A key reason these late-stage conversations become flashpoints for discrimination is the dilution of older social markers—like surnames or accents—in today's interconnected labor markets. As people move more freely and higher education expands, it becomes harder to "read" someone's background from a résumé alone. Ironically, this pushes status-based screening into deeper, less structured interactions—namely personal interviews—where intangible signals resurface and employers lean on informal dialogues that often reproduce entrenched privilege. Thus, the weakening of obvious status cues can paradoxically magnify the impact of subjective impressions at the moment they're most harmful.

The findings also shed new light on how affirmative action gains in education do not necessarily translate into equitable labor-market outcomes. Even after disadvantaged groups surmount formidable academic hurdles—often aided by affirmative action policies—they might encounter hiring practices that penalize the very backgrounds those policies aimed to uplift. By highlighting how last-stage assessments can erode earlier successes, this research underscores a broader pattern: organizations may overtly endorse inclusion yet retain practices that let structural inequalities reassert themselves in final hiring decisions. The resulting inefficiency is twofold: on one hand, firms lose out on top performers who, as promotion data suggest, excel when given a fair chance; on the other, historically marginalized groups bear disproportionate burdens just as they approach the threshold of professional mobility.

Beyond any single hiring decision, these dynamics reflect broader stakes for organizations and their talent strategies. When final-round "fit" assessments systematically exclude underrepresented talent, firms lose a pipeline of potential star employees. Hiring managers calling the shots often prioritize short-term rapport or "chemistry," sidestepping more robust measures of future productivity (Benson and Shaw, 2024; Hoffman and Stanton, 2024). By pinpointing the disconnect between late-stage interview impressions and subsequent on-the-job success, this study offers clear evidence that rethinking "fit" is not just an equity imperative but also a competitive necessity. Rather than eliminating interviews, companies can acknowledge these subtler, late-stage vulnerabilities and mitigate the pitfalls through standardized questioning, stronger post-hire accountability for

evaluators, training to recognize hidden prejudices, or oversight mechanisms that curb unilateral gatekeeper power.

Moreover, this paper's findings illustrate how subjective conversations can disadvantage not only caste groups but also other communities with concealable identities—whether LGBTQ+, differently abled, or from low-income backgrounds that don't announce themselves on a résumé. Such identities often surface only when candidates speak at length about their experiences or personal histories. Policy discussions surrounding workforce diversity could thus benefit from step-by-step audits of each hiring phase. Rather than focusing solely on earlier, more straightforward forms of prejudice (e.g., résumé callbacks), effective reforms must also address the concealed biases that surface once more personal information is revealed.

Looking ahead, as remote work expands and algorithmic tools reshape hiring, organizations must decide whether to preserve traditional "chemistry"-driven interviews or adopt more transparent, skill-oriented approaches. While AI-driven screenings and remote formats may mitigate some biases, they can also magnify status cues embedded in candidates' backgrounds or online footprints (Albright, 2024). Balancing managerial discretion with structured evaluation protocols—such as standardized behavioral interviews—could foster both equity and more accurate talent identification. Ultimately, this choice will influence not only firms' competitiveness but also the trajectory of social mobility in elite labor markets worldwide.

References

[1] Abaluck, J., L. Agha, C. Kabrhel, A. Raja, and A. Venkatesh (2016). The determinants of productivity in medical testing: Intensity and allocation. *American Economic Review 106*(12), 3730–3764.

[2] Albright, A. (2024). The hidden effects of algorithmic recommendations. Working paper, Federal Reserve Bank of Minneapolis.

[3] Altonji, J. G. and R. M. Blank (1999). *Race and Gender in the Labor Market. Handbook of Labor Economics*, ed. by O. C. Ashenfelter and D. Card, Elsevier, vol. 3C, chap. 48.

[4] Asher, S., P. Novosad, and C. Rafkin (2024). Intergenerational mobility in india: New measures and estimates across time and social groups. *AEJ: Applied Economics 16*(2), 66–98.

[5] Bayly, S. (2008). *Caste, Society and Politics in India from the Eighteenth Century to the Modern Age*. Chicago: Cambridge University Press.

[6] Benson, A., S. Board, and M. Meyer-ter-Vehn (2024). Discrimination in hiring: Evidence from retail. *Review of Economic Studies 91*(4), 1956–1987.

[7] Benson, A., D. Li, and K. Shue (2019). Promotions and the peter principle. *The Quarterly Journal of Economics 134*(4), 2085–2134.

[8] Benson, A., D. Li, and K. Shue (2023). "Potential" and the gender promotion gap. R&R, *American Economic Review*.

[9] Benson, A. and K. Shaw (2024). What do managers do? *Annual Review of Economics*. Forthcoming.

[10] Beteille, A. (1965). *Caste, Class and Power: Changing Patterns of Stratification in a Tanjore Village*. University of California Press.

[11] Beteille, A. (1969). *Castes: Old and New, Essays in Social Structure and Social Stratification*. Asia Publishing House.

[12] Bhavnani, R. R. and A. Lee (2021). Does affirmative action worsen bureaucratic performance? evidence from the indian administrative service. *American Journal of Political Science 65*(1), 5–20.

[13] Bloom, N., R. Sadun, and J. Van Reenan (2012). The organization of firms across countries. *The Quarterly Journal of Economics 127*(4), 1663–1705.

[14] Bloom, N., R. Sadun, and J. Van Reenen (2010). Recent advances in the empirics of organizational economics. *Annual Review of Economics 2*, 105–137.

[15] Bloom, N., R. Sadun, and J. Van Reenen (2017). Management as a technology? NBER Working Paper 22327.

[16] Bloom, N. and J. Van Reenen (2007). Measuring and explaining management practices across firms and countries. *The Quarterly Journal of Economics 122*(4), 1351–1408.

[17] Bourdieu, P. (1986). The forms of capital. In J. G. Richardson (Ed.), *Handbook of Theory and Research for the Sociology of Education*, pp. 241–258. Greenwood Press.

[18] BusinessLine (2018). An Inclusive India Inc. Business Line, January 13, 2018.

[19] California v. Cisco Systems, Inc. (2022). Dept. of Fair Employment and Housing v. Cisco Systems, Inc. Docket No. H048962.

[20] Cappelli, P. (2019). Your approach to hiring is all wrong. *Harvard Business Review*. Accessed: 2024-08-15.

[21] Chandra, A. and D. O. Staiger (2010). Identifying provider prejudice in healthcare. Technical Report 16382, National Bureau of Economic Research.

[22] Chandra, A. and D. O. Staiger (2020). Identifying sources of inefficiency in health care. *Quarterly Journal of Economics 135*(2), 785–843.

[23] Chetty, R., D. J. Deming, and J. N. Friedman (2023). Diversifying society's leaders? the determinants and causal effects of admission to highly selective private colleges. NBER Working Paper 31492.

[24] Clark, G. (2015). *The Son Also Rises: Surnames and the History of Social Mobility*. The Princeton Economic History of the Western World.

[25] Datta, S. (2017). *A History of the Indian University System: Emerging from the Shadows of the Past*. Palgrave Macmillan.

[26] De Janvry, A., G. He, E. Sadoulet, S. Wang, and Q. Zhang (2023). Subjective performance evaluation, influence activities, and bureaucratic work behavior: Evidence from china. *American Economic Review 113*(3), 766–799.

[27] Deshpande, A. (2011). *The Grammar of Caste: Economic Discrimination in Contemporary India*. Oxford: Oxford University Press.

[28] Durlauf, S. N. (1996). A theory of persistent income inequality. *Journal of Economic Growth 4*(1), 75–93.

[29] Durlauf, S. N. (2004). Neighborhood effects. In J. V. Henderson and J.-F. Thisse (Eds.), *Handbook of Regional and Urban Economics*, Volume 4, pp. 2173–2242. Elsevier.

[30] Durlauf, S. N. and M. Fafchamps (2006). Social capital. In P. Aghion and S. N. Durlauf (Eds.), *Handbook of Economic Growth*. Amsterdam: North Holland.

[31] Durlauf, S. N. and Y. Ioannides (2010). Social interactions. *Annual Review of Economics 2*, 451–478.

[32] Eames, T. (2024). Taryn versus taryn (she/her) versus taryn (they/them): A field experiment on pronoun disclosure and hiring discrimination. Working Paper.

[33] Eames, T. (2025). How does sex-based occupational discrimination affect nonbinary applicants? *AEA Papers and Proceedings*.

[34] Economic Advisory Council to the Prime Minister (2022). The State of Inequality in India Report. Government of India.

[35] Fenizia, A. (2022). Managers and productivity in the public sector. *Econometrica 90*(3), 1063–1084.

[36] Fisman, R., J. Shi, Y. Wang, and R. Xu (2018). Social ties and favoritism in chinese science. *Journal of Political Economy 126*(3), 1134–1171.

[37] Folke, O. and J. Rickne (2016). The glass ceiling in politics: Formalization and empirical tests. *Comparative Political Studies 49*(5), 567–599.

[38] Friebel, G., M. Heinz, M. Krueger, and N. Zubanov (2022). Team incentives and performance: Evidence from a retail chain. *American Economic Review 107*(8), 2168–2203.

[39] Friebel, G. and M. Raith (2024). Talent management: The role of bosses. R&R, *Journal of Labor Economics*.

[40] Friedman, S. and D. Laurison (2020). *The Class Ceiling: Why it Pays to be Privileged?* Bristol University Press, Policy Press.

[41] Gibbons, R. and J. Roberts (2013). Economic theories of incentives in organizations. *The Handbook of Organizational Economics*.

[42] Hoffman, M., L. B. Kahn, and D. Li (2018). Discretion in Hiring. *The Quarterly Journal of Economics 133*(2), 765–800.

[43] Hoffman, M. and C. T. Stanton (2024). People, practices, and productivity: A review of new advances in personnel economics. NBER Working Paper 32849.

[44] Hoffman, M. and S. Tadelis (2021). People management skills, employee attrition, and manager rewards: An empirical analysis. *Journal of Political Economy 129*(1), 243–285.

[45] Jodhka, S. S. and J. Naudet (2019). *Mapping the Elite: Power, Privilege, and Inequality*. Oxford University Press.

[46] Jodhka, S. S. and K. Newman (2007). In the name of globalisation: Meritocracy, productivity and the hidden language of caste. *Economic and Political Weekly 42*(41), 4125–4132.

[47] Jodkha, S. S. (2017). *Caste in Contemporary India*. Routledge India.

[48] Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

[49] Kaila, M., E. Nix, and K. Riukula (2024). The impact of an early career shock on intergenerational mobility. *Journal of Labor Economics*. Forthcoming.

[50] Kline, P., E. Rose, and C. Walters (2022). Systemic discrimination among large u.s. employers. *Quarterly Journal of Economics 137*(4), 1963–2036.

[51] Knobe, J. and S. Nichols (2008). An introduction to experimental philosophy. In *Experimental Philosophy*, pp. 3–34. Oxford: Oxford University Press.

[52] Kohlhepp, J. (2024). The inner beauty of firms. R&R, *Journal of Political Economy*.

[53] Lamont, M. (1992). *Money, Morals, and Manners: The Culture of the French and the American Upper-Middle Class*. Chicago: University of Chicago Press.

[54] Lamont, M. (2009). *How Professors Think: Inside the Curious World of Academic Judgment*. Cambridge, MA: Harvard University Press.

[55] Lamont, M. and A. Lareau (1988). Cultural capital: Allusions, gaps and glissandos in recent theoretical developments. *Sociological Theory 6*(2), 153–168.

[56] Laurison, D. and S. Friedman (2019). *The Class Ceiling: Why it Pays to be Privileged*. Bristol, UK: Policy Press.

[57] Lazear, E. P. (2009). *Personnel Economics*. Cambridge, MA: MIT Press.

[58] Lazear, E. P. and P. Oyer (2012). Personnel economics. In R. Gibbons and J. Roberts (Eds.), *The Handbook of Organizational Economics*, pp. 479–519. Princeton University Press.

[59] Macdonald, D., J. Montonen, and E. Nix (2024). Dating and breaking up with the boss: Benefits, costs, and spillovers. Working Paper.

[60] Madheswaran, S. (2008). Is affirmative action policy for private sector necessary? *Indian Journal of Industrial Relations 44*(2), 164–184.

[61] Ministry of Law and Justice, Government of India (2021). The constitution of india. Article 15 explicitly prohibits discrimination on grounds of religion, race, caste, sex, or place of birth.

[62] Minni, V. (2023). Making the invisible hand visible: Managers and the allocation of workers to jobs. R&R, *Quarterly Journal of Economics*.

[63] Minni, V. (2024). Global managers, local workers: Wage setting inside a multinational firm. *AEA Papers & Proceedings 114*, 586–91.

[64] Mishra, N. (2015). India and colorism: The finer nuances. *Washington University Global Studies Law Review 14*, 725–750.

[65] Newman, K. and S. Thorat (2012). *Blocked by Caste: Economic Discrimination in Modern India*. Oxford University Press.

[66] Novosad, P., S. Asher, C. Farquharson, and E. Iljazi (2024). Access to opportunity in the sciences: Evidence from the nobel laureates. Working Paper.

[67] Obermeyer, Z., B. Powers, C. Vogeli, and S. Mullainathan (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science 366*(6464), 447–453.

[68] Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics 37*(2), 187–204.

[69] Oyer, P. and S. Schaefer (2011). Personnel economics: Hiring and incentives. *Handbook of Labor Economics 4b*, 1769–1823.

[70] Parameswaran, R. and K. Cardoza (2015). Melanin on the margins: Advertising and the cultural politics of fair/light/white beauty in india. *Journalism & Communication Monographs 11*(3), 213–274.

[71] Pinker, S. (2021). *Rationality: What It Is, Why It Seems Scarce, Why It Matters*. New York: Viking.

[72] Prendergast, C. (2007). The motivation and bias of bureaucrats. *The American Economic Review 97*(1), 180–196.

[73] Rivera, L. A. (2012). Hiring as cultural matching: The case of elite professional service firms. *American Sociological Review 77*(6), 999–1022.

[74] Rivera, L. A. (2015). *Pedigree: How Elite Students Get Elite Jobs*. Princeton University Press.

[75] Rivera, L. A. (2020). Employer decision making. *Annual Review of Sociology 46*(1), 215–232.

[76] Rose, E. K. (2023, September). A constructivist perspective on empirical discrimination research. *Journal of Economic Literature 61*(3), 906–923.

[77] Sacerdote, B., D. O. Staiger, and M. Tine (2025, January). How test optional policies in college admissions disproportionately harm high achieving applicants from disadvantaged backgrounds. Working Paper 33389, National Bureau of Economic Research. Available at http://www.nber.org/papers/w33389.

[78] Sarsons, H. (2017). Recognition for group work: Gender differences in academia. *AEA Papers and Proceedings 107*(5), 141–145.

[79] Sarsons, H., K. Gërxhani, E. Reuben, and A. Schram (2021). Gender differences in recognition for group work. *Journal of Political Economy 129*(1), 101–147.

[80] Sarsons, H., J. Hjort, and X. Li (2022). Across-Country Wage Compression in Multinationals. R&R, *American Economic Review*.

[81] Stansbury, A. and K. Rodriguez (2024). The class gap in career progression: Evidence from academia. R&R, *Econometrica*.

[82] Stansbury, A. and R. Schultz (2023). The economics profession's socioeconomic diversity problem. *Journal of Economic Perspectives 37*(4), 207–230.

[83] Subramanian, S. (2019). India's elite colleges and the employment divide. *Economic and Political Weekly 54*(2), 42–48.

[84] Todorov, A. (2017). *Face Value: The Irresistible Influence of First Impressions*. Princeton, NJ: Princeton University Press.

[85] Wiltshire, C. R. (2020). *Uniformity and Variability in the Indian English Accent*. Cambridge University Press.

[86] Yenmez, M. B. and T. Sönmez (2022). Affirmative action in India via vertical, horizontal, and overlapping reservations. *Econometrica 90*(3), 1143–1176.

A ONLINE APPENDIX

A.1 Representativeness of the Job Fair of the Elite College

This section argues that the on-campus job fair of the elite college provides a representative view of how elite college graduates in India transition into elite jobs. This is because of the following four reasons:

1. **Elite colleges have similar on-campus job fairs.** Elite Indian colleges have similar on-campus job fairs due to an extended process of historical imitation. Post independence, the earliest elite Indian colleges were built in the 1950s and were closely modeled on elite US universities, particularly MIT and Stanford. The earliest elite Indian colleges served as role models for elite Indian colleges built later, which imitated key aspects such as academic calendars, faculty-to-student ratios, and job fairs. Thus, almost all elite Indian colleges today have similar mechanisms for selectively placing graduates into elite jobs (Central Placement Cell, Delhi University). Examples of elite Indian colleges sharing common job fairs include 23 IITs, 20 IIMs, and several colleges under the umbrella of the prestigious Delhi University.

2. **Elite college graduates primarily work in elite entry-level jobs.** These jobs are primarily situated in the private sector and in major multinational corporations. See Subramanian (2019) and Jodhka and Naudet (2019) for a comprehensive overview.

3. **Elite entry-level jobs almost exclusively hire from elite colleges.** Firms in my sample visit about 40 other Indian colleges every year, each with over 1000 students.[OA.1] See Subramanian (2019) and Jodhka and Naudet (2019) for a comprehensive overview.

4. **Small proportions of students skip the job fairs of elite colleges.** Students must register for the formal on-campus job fairs of elite colleges, but otherwise must "deregister" or formally drop out. Only a small proportion (less than 5%) of college graduates in my setting skip the formal placement process for full-time jobs (Section 3.2). Given the similarity of the placement process across elite Indian colleges, this proportion is likely to be similar in other institutions as well. See Subramanian (2019) for more details.

   The discussion above implies that the job placement process studied in this paper offers a representative window into how elite college graduates transition into elite entry-level jobs in the Indian private sector.

[OA.1] These details were calculated by scraping the identities of firms from the placement websites of many Indian colleges. More details available upon request.

A.2 More Details on Timeline and Screening Rounds

This section provides more details on the job search timeline and data from each of the four screening rounds conducted by each job.

Four Screening Rounds Per Job. Each job in my sample conducts four screening rounds and eliminates students at each of them. The four job screening rounds are as follows:

1. **Application reading.** In the first screening round, employers read applications and determine whether to advance students further. This "first cut" is typically made using a combination of college GPA, major, degree, and other characteristics on student CVs (e.g., summer internships and exceptional performance in international competitions). Informal conversations with employers suggest that college GPA is one of the main filters used during the application reading round. The first screening round is usually conducted in September after applications are submitted.

   Following the application reading round, the second and third screening rounds are conducted between September to late November.

2. **Written aptitude tests.** The second screening round is a written aptitude test. These tests, also called "technicals," are mostly conducted online. Technicals differ by job roles and usually include coding tests for manufacturing or technology jobs and case-study-based tests ("general aptitude tests") for consulting jobs. Written tests are almost always conducted online.

   It is typical for a firm advertising many job profiles to conduct common written tests for all advertised profiles within that firm. For example, if Nvidia advertises 5 job profiles in a given year, it might choose to conduct a common written test ("technical") for all profiles at once.

3. **Group debates.** The third screening round is a "group debate." It is a *non-technical* round. Group debates are conducted for all jobs in my sample, whether client-facing or non-client-facing.

   a. **Group debates are conducted in multiple batches per job, with each batch evaluated by different members of the full recruiting team.** Each batch comprises a total of 30-35 students. This batch is further divided into two teams of 15-20 students each. Group debates last a few hours and assess a wide array of socio-emotional skills, including: communication, mannerisms, consensus building, collegiality, confidence, and teamwork. Group debate topics are general in scope (e.g., "Is social media a boon or a bane?").

      During group debates, recruiters make students discuss a general or non-technical topic with one team speaking for each side. Meanwhile, recruiters behave like passive observers, with their only active role consisting of duties

such as starting and ending the discussion on time and ensuring decorum. Discussion moderators organically emerge from the group of students participating in the debate.

b. **Many firms conduct common group debates for multiple job positions.** Like written tests, it is typical for a firm advertising many job profiles to conduct common group debates for all advertised profiles within that firm. For example, if Google advertises 5 job profiles in a given year, it might choose to conduct a common group debate for all profiles at once. This practice reflects the fact that employers care about basic soft skills and the ability to work in teams across job profiles in different core roles within a firm (e.g., software engineer or program manager at Google).

4. **Personal interview.** The fourth and final screening round is personal interview. Personal interviews are usually conducted between December and January.

   a. **Personal interviews are conducted in multiple batches per job, with each batch evaluated by different members of the full recruiting team.** Like group debates, personal interviews are typically conducted in multiple batches per job, with each batch evaluated by different members of the full recruiting team.

   b. **Many firms conduct common interviews for multiple job positions.** Like written tests and group debates, firms typically conduct common personal interviews for multiple job positions. For example, Uber might be hiring a program manager, data scientist, a frontend engineer, and a backend engineer in a given year. Interviews for all these job positions within Uber are typically bundled into one, again reflecting the fact that many common skills or aspects related to candidate background or "fit" are relevant in the employer's eye across all job positions within a firm.

   c. **Almost all interviewers are likely Indians.** Recall, I do not directly have information on interviewer nationality. However, since more than 98% of the jobs are located in India, it is reasonable to assume that recruitment is typically conducted by Indian hiring professionals or managers (Section 4.5).

A.3 Extending the Simple Model of Late-Stage Discrimination with Imperfect Information

In Section 5.6.1, we assumed that employers observe each candidate's true screening-round performance perfectly, i.e., $M_{s,i,g} = A_{i,g}$ at round s. We now show that *all* of the key insights—particularly how job-selectivity variation identifies discrimination versus ability deficits—remain valid even when employers observe *noisy* performance signals.

Setup. As in Section 5.6.1, there are two groups (castes) $g \in \{0, 1\}$, a continuum of jobs $j \in \mathcal{J}$, and up to four sequential screening rounds $\{1, 2, 3, 4\}$. Rounds $1, 2, 3$ constitute pre-interview filtering (e.g., applications, written tests, group debates), while round 4 is the final personal interview. Each candidate i from group g has a *latent ability* $A_{i,g}$. Passing round s requires

$$E\big[M_{s,i,g} \,\big|\mathcal{I}_{s-1}\big] \;\geq\; \tau_{s,j}, \tag{OA.1}$$

where $\mathcal{I}_{s-1}$ is the employer's information after observing signals in rounds $1, \ldots, s-1$. As before, each round's threshold $\tau_{s,j}$ is weakly increasing with s and may differ across jobs j.

Noisy Signals per Round. At each stage s, the employer observes

$$M_{s,i,g} \;=\; A_{i,g} \;+\; \eta_{s,i,g},$$

where $\eta_{s,i,g}$ is a zero-mean noise term. We relax the assumption of independence across candidates and rounds, requiring only that:

(i) $\mathrm{E}[\eta_{s,i,g}] = 0$ for all s, i, g, and

(ii) $\eta_{s,i,1} \overset{d}{=} \eta_{s,i,0}$, so that for any subset of candidates (including those who have passed earlier rounds), there is no group-level shift or difference in noise distributions that would systematically favor one caste.

We allow that $(\eta_{s,i,g})$ may be correlated across s or i (e.g., a candidate's performance in different rounds might be interdependent), and may also depend on the conditioning event of "having passed earlier rounds." The key point is that *noise does not differ across castes in a way that would introduce persistent group-level biases.* Under these assumptions, any stable group differences cannot be attributed to noise alone.

Updating Beliefs and Passing Rounds. At each stage s, the employer forms an updated belief about $A_{i,g}$ given $\mathcal{I}_s \equiv \{M_{1,i,g}, \ldots, M_{s,i,g}\}$. Let

$$E[A_{i,g} \mid \mathcal{I}_s] \quad \text{and} \quad Var(A_{i,g} \mid \mathcal{I}_s)$$

denote the posterior mean and variance of $A_{i,g}$. Intuitively, with more signals $\{M_{1,i,g}, \ldots, M_{s,i,g}\}$, beliefs about each candidate's ability become more precise, and the passing condition (OA.1) can be evaluated in terms of $E[M_{s,i,g} \mid \mathcal{I}_{s-1}]$.

Final Interview Thresholds and Caste Revelation. As before, let $\tau_{4,j}^{(g)}$ be the *group-specific* final-round threshold at job j. If caste becomes salient at the interview for group 1, a higher threshold ($\tau_{4,j}^{(1)} > \tau_{4,j}^{(0)}$) would imply discrimination. Our goal is to see how variation in $\tau_{4,j}$ across jobs still allows us to disentangle true ability differences from discrimination, even when $M_{s,i,g}$ include noise.

A.3.1 No Early Caste Penalty Implies No Large Right-Tail Deficit Before Interviews

If there were a substantial right-tail *ability* deficit for group 1, we would see an especially large gap in pass rates at earlier stages for *highly selective* jobs (i.e., with large $\tau_{s,j}$). Formally:

Proposition 3 (No Early Gap $\implies$ No Right-Tail Deficit Under Imperfect Information)**.** Consider any pre-interview stage $s < 4$. If group 1 truly has fewer high-ability candidates in the top tail, then for sufficiently large $\tau_{s,j}$,

$$P\big(\text{Pass}_{s,i,j} \mid g = 1, \mathcal{I}_{s-1}\big) \;<\; P\big(\text{Pass}_{s,i,j} \mid g = 0, \mathcal{I}_{s-1}\big).$$

Hence, if empirically no such gap arises—even in the most selective jobs—we can rule out a meaningful right-tail deficit for group 1 at the earlier stages.

Sketch. With noisy signals, passing round s depends on the posterior mean $E[M_{s,i,g} \mid \mathcal{I}_{s-1}]$. For large $\tau_{s,j}$, the employer needs $E[A_{i,g} \mid \mathcal{I}_{s-1}]$ to be sufficiently high to clear that threshold. If group 1 truly had fewer high-ability individuals, then as $\tau_{s,j}$ grows, the fraction of group 1 candidates surpassing it falls *faster* than for group 0. This mirrors the perfect-information argument in Proposition 1 and further relies on the *regularly decaying tail* assumption (e.g., $P(M_{s,i,g} \geq m) \sim C_g\, m^{-\alpha}$ for large m), which ensures that higher cutoffs intensify right-tail gaps. Because noise is symmetric across castes, it cannot mask or invert these latent-ability differences at sufficiently high cutoffs. Hence, if we observe no substantial disadvantage for group 1 at early rounds even in highly selective jobs, we can rule out a meaningful right-tail deficit for that group. □

A.3.2 Stable Promotion Gaps Across Selectivity Still Imply Discrimination

Just as in Section 5.6.1, consider post-hire performance Y_i tied to underlying ability $A_{i,g}$. A candidate is promoted if $Y_i \geq \kappa_j$. Let final-round passing require $E[M_{4,i,g} \mid \mathcal{I}_3] \geq \tau_{4,j}^{(g)}$. We then define a *promotion gap* across selectivity tiers. Suppose jobs in the top $(1-q)$ fraction have cutoffs $\tau_{4,j}^{q}$. If both castes face $\tau_{4,j}^{(0)} = \tau_{4,j}^{(1)} = \tau_{4,j}^{q}$, then as $\tau_{4,j}^{q} \to \infty$, average hired performance converges to the *common marginal candidate*, making the promotion gap shrink to zero. By contrast, if $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} + \delta$ ($\delta > 0$), group 1's *marginal* hires have strictly higher *posterior* means, yielding a stable advantage in realized performance.

Proposition 4 (Stable Promotion Gaps Under Imperfect Information Identify Discrimination)**.** Under the same regularly decaying tail assumption used in Proposition 2 (e.g. for large m, $P(M_{4,i,g} \geq m) \sim C_g\, m^{-\alpha}$ with $\alpha > 1$), if the final-round threshold for group 1 exceeds that for group 0 by a constant $\delta > 0$, then the average hired performance for group 1 remains approximately δ units higher than group 0, even in highly selective jobs. Consequently, if promotion-rate gaps do not vanish at higher cutoffs, we can reject an

unobserved-ability explanation in favor of discrimination.

Sketch. The argument parallels Proposition 2. Under imperfect information, stage-4 passage requires $E[M_{4,i,g} \mid \mathcal{I}_3] \geq \tau_{4,j}^{(g)}$. If group 1 faces a higher threshold $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} + \delta$, then a typical group-1 hire must satisfy

$$E[A_{i,1} \mid \mathcal{I}_3] \;\approx\; \tau_{4,j}^{(0)} + \delta \;-\; E[\eta_{4,i,1} \mid \mathcal{I}_3].$$

Because noise is identically distributed across groups (mean-zero, no group-level bias), raising group 1's threshold by δ effectively *shifts* group 1's final hires to a higher region of the ability distribution.

As $\tau_{4,j}$ rises across selectivity tiers, the average performance of each group's hired candidates converges to the respective threshold they must clear. Since group 1's threshold is δ higher, group 1's average hired performance remains about δ above group 0's. By contrast, if both groups share the same threshold, they converge to the same marginal point, so any promotion gap collapses to zero. The regularly decaying tail assumption ensures that these marginal hires dominate average outcomes at high cutoffs, preserving a stable δ-sized gap under discrimination. □

Key Takeaway. Allowing for imperfect information changes the exact pass mechanics but preserves the same identifying variations. If there is *no* sign of a pre-interview deficit at high cutoffs and we see a stable or rising performance gap post-hire, the best explanation remains that group 1 faces a stricter interview threshold once caste is revealed, rather than any hidden ability shortfall.

A.4 Tables and Figures

Table OA.1: **Caste Differences in Mean Salary Among Graduates**

	Dependent variable:			
	log (Salary)			
	Without Controls	With Controls	Client Facing	Non-Client Facing
	(1)	(2)	(3)	(4)
Disadv. Caste	-0.231^{***}	-0.103^{***}	-0.107^{***}	-0.089^{***}
	(0.020)	(0.017)	(0.027)	(0.022)
Disadv. Caste (Oster)		-0.076^{***}	-0.095^{***}	-0.072^{***}
		(0.019)	(0.026)	(0.022)
Constant	11.235^{***}	10.863^{***}	10.883^{***}	10.848^{***}
	(0.012)	(0.040)	(0.057)	(0.052)
Observations	2,527	2,527	849	1,678
Adjusted R^2	0.049	0.412	0.549	0.382
Year FE	✓	✓	✓	✓
Other Controls	✓	✓	✓	✓
Job FE				

Notes: Online Appendix Table OA.1 includes estimates from an earnings regression run on the sample of all students who graduated with jobs obtained through the college's job fair for full-time jobs. The dependent variable is log earnings. The regressions include detailed controls including measures of pre-college skills, college academic performance, previous labor market experience, and other employer-relevant skills such as major and degree. Column (1) reports estimates from a linear earnings regression without controls. Column (2) reports estimates from a linear earnings regression with controls. Columns (3) and (4) report estimates from a linear earnings regression with controls run separately for client-facing and non-client-facing jobs, respectively. Note that the earnings regressions do not include job fixed effects because there are no caste differences in pay for identical jobs. Bias-adjusted coefficients are calculated using Oster (2019). Following Oster's recommended benchmarks, I set $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$, implying that unobservables are assumed to have the same degree of selection as observables, and the total explanatory power can increase by up to 30% beyond the observed R^2. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table OA.2: **Caste Differences in GPA Between Those Omitted From the Earnings Regression Versus Those Who Are Not**

	Dependent variable:	
	log (GPA)	
	(1)	(2)
Disadv. Caste	−0.033***	−0.032***
	(0.004)	(0.004)
Omitted	−0.167***	−0.158***
	(0.005)	(0.005)
Omitted * Disadv. Caste	−0.071***	−0.070***
	(0.007)	(0.007)
Constant	2.147***	2.127***
	(0.004)	(0.008)
Observations	4,164	4,164
Adjusted R^2	0.546	0.559
Year FE	✓	✓
Major		✓
Degree		✓
Job FE		
Note:	*p<0.1; **p<0.05; ***p<0.01	

Notes: Online Appendix Table OA.2 shows the caste differences in GPA between those omitted from the earnings regression versus those who are not. The "omitted" category includes those who deregistered from the placement process and those who registered but did not receive any job offers (Section 5.2). The latter group often participates in the aftermarket and usually pursue higher education or entrepreneurship (Section 3.2). Column (1) only includes year fixed effects. Column (2) adds major and degree as controls. Note that the regressions do not have job fixed effects because these can typically only be included for jobs in my main sample. Therefore, these job fixed effects cannot be separately identified from the coefficient on "omitted."

Table OA.3: **Caste Differences in GPA At Each Job Search Stage**

	Dependent variable:					
	log (GPA)					
	App	Test	Debate	Interview	Offer	Choice
	(1)	(2)	(3)	(4)	(5)	(6)
Disadv. Caste	−0.111***	−0.049***	−0.039***	−0.019***	0.010***	0.009***
	(0.005)	(0.003)	(0.002)	(0.001)	(0.002)	(0.002)
Disadv. Caste (Oster)	−0.104***	−0.035***	−0.034***	−0.014***	0.015***	0.014***
	(0.006)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)
Constant	2.057***	2.131***	2.163***	2.212***	2.214***	2.199***
	(0.028)	(0.024)	(0.023)	(0.006)	(0.006)	(0.028)
Mean GPA	7.776	7.967	8.17	8.26	8.419	8.419
Observations	4,005	3,618	3,266	3,031	2,527	2,527
Adjusted R^2	0.264	0.726	0.745	0.801	0.822	0.822
Year FE	✓	✓	✓	✓	✓	✓
Job FE	✓	✓	✓	✓	✓	✓
Major	✓	✓	✓	✓	✓	✓
Degree	✓	✓	✓	✓	✓	✓
Note:						*p<0.1; **p<0.05; ***p<0.01

Notes: Online Appendix Table OA.3 compares the *between*-caste and *within*-job GPA differences at each job search stage. These differences are *controlled*, with the linear regressions including standard controls such as major, degree, and year, as well as job fixed effects. Bias-adjusted coefficients are calculated using Oster (2019). Following Oster's recommended benchmarks, I set $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$, implying that unobservables are assumed to have the same degree of selection as observables, and the total explanatory power can increase by up to 30% beyond the observed R^2. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table OA.4: **Caste Differences in Mean Salary at Each Job Search Stage in the Full Sample**

	Dependent variable:					
	log (Salary)					
	App	Test	Debate	Interview	Offer	Choice
	(1)	(2)	(3)	(4)	(5)	(6)
Disadv. Caste	0.004	−0.009	−0.006	−0.018*	−0.105***	−0.103***
	(0.005)	(0.007)	(0.007)	(0.010)	(0.013)	(0.017)
Disadv. Caste (Oster)	0.005	0.010	0.011	0.012	−0.079***	−0.076***
	(0.005)	(0.008)	(0.008)	(0.011)	(0.014)	(0.019)
Constant	11.032***	10.904***	10.848***	10.758***	10.672***	10.863***
	(0.011)	(0.015)	(0.015)	(0.023)	(0.030)	(0.040)
Observations	4,005	3,618	3,266	3,031	2,527	2,527
Adjusted R^2	0.578	0.491	0.547	0.549	0.488	0.412
Year FE	✓	✓	✓	✓	✓	✓
Other Controls	✓	✓	✓	✓	✓	✓
Job FE						

Notes: Online Appendix Table OA.4 includes estimates from an earnings regression run on the sample of all students who remain in contention for some job at a given job search stage. For example, column (1) includes the sample of all students who apply for jobs recruiting from the college's job fair. Similarly, column (2) includes the sample of all students who remain in contention for some job at the written test round, and so on. The dependent variable is log earnings. I include detailed controls, including measures of pre-college skills, college academic performance, previous labor market experience, and other employer-relevant skills such as major and degree. Each column is a separate regression and includes these controls. Note that the earnings regressions do not include job fixed effects because there are no caste differences in pay for identical jobs. Bias-adjusted coefficients are calculated using Oster (2019). Following Oster's recommended benchmarks, I set $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$, implying that unobservables are assumed to have the same degree of selection as observables, and the total explanatory power can increase by up to 30% beyond the observed R^2. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table OA.5: **Caste Differences in the Success Rate at Each Job Search Stage Among Client-Facing Jobs**

	Success App (1)	Success Test (2)	Success Debate (3)	Success Interview (4)
Disadv. Caste (OLS)	-0.0066***	-0.0007	-0.0136***	-0.3062***
	(0.0022)	(0.0030)	(0.0046)	(0.0081)
Disadv. Caste (Oster)	0.002	0.006	-0.015***	-0.314***
	(0.002)	(0.004)	(0.005)	(0.009)
Mean Success Rate	0.5770	0.7860	0.2590	0.2600
Observations	117,770	67,927	53,422	13,842
Adjusted R^2	0.61588	0.37467	0.01189	0.24298
Year FE	✓	✓	✓	✓
Job FE	✓	✓	✓	✓
Other Controls	✓	✓	✓	✓

Notes: Online Appendix Table OA.5 shows the caste difference in the success rate at each stage of job search among client-facing jobs. Each observation is a student who remains in contention at a given screening round in a given job in a given year. Each column reports a separate linear probability regression and includes controls. These regressions also include job fixed effects. Note that some firms conduct the same screening round (written tests, group debates, interviews) for multiple job positions. However, since these job positions within the same firm represent different jobs, the regressions consider these jobs and the corresponding advance-or-out decisions made by them separately. Bias-adjusted coefficients are calculated using Oster (2019). Following Oster's recommended benchmarks, I set $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$, implying that unobservables are assumed to have the same degree of selection as observables, and the total explanatory power can increase by up to 30% beyond the observed R^2. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

Table OA.6: **Caste Differences in the Success Rate at Each Job Search Stage Among Non-Client-Facing Jobs**

	Success App (1)	Success Test (2)	Success Debate (3)	Success Interview (4)
Disadv. Caste (OLS)	0.0024	0.0011	0.0070**	-0.2867***
	(0.0018)	(0.0022)	(0.0034)	(0.0060)
Disadv. Caste (Oster)	0.008***	0.014***	0.007**	-0.286***
	(0.002)	(0.002)	(0.003)	(0.007)
Mean Success Rate	0.6140	0.8250	0.2580	0.2620
Observations	187,771	115,223	95,067	24,536
Adjusted R^2	0.60783	0.32634	0.02270	0.24723
Year FE	✓	✓	✓	✓
Job FE	✓	✓	✓	✓
Other Controls	✓	✓	✓	✓

Notes: Online Appendix Table OA.6 shows the caste difference in the success rate at each stage of job search among non-client-facing jobs. Each observation is a student who remains in contention at a given screening round in a given job in a given year. Each column reports a separate linear probability regression and includes controls. These regressions also include job fixed effects. Note that some firms conduct the same screening round (written tests, group debates, interviews) for multiple job positions. However, since these job positions within the same firm represent different jobs, the regressions consider these jobs and the corresponding advance-or-out decisions made by them separately. Bias-adjusted coefficients are calculated using Oster (2019). Following Oster's recommended benchmarks, I set $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$, implying that unobservables are assumed to have the same degree of selection as observables, and the total explanatory power can increase by up to 30% beyond the observed R^2. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table OA.7: **Caste Differences in Mean Salary at Each Job Search Stage in the Identifiable Sample**

	Dependent variable:					
	log (Salary)					
	App	Test	Debate	Interview	Offer	Choice
	(1)	(2)	(3)	(4)	(5)	(6)
Disadv. Caste (OLS)	0.012	-0.068^{***}	-0.034^{*}	-0.105^{***}	-0.146^{***}	-0.107^{**}
	(0.014)	(0.019)	(0.018)	(0.027)	(0.033)	(0.049)
Disadv. Caste (Oster)	0.013	-0.049^{***}	-0.016	-0.078^{**}	-0.119^{***}	-0.082
	(0.016)	(0.021)	(0.022)	(0.032)	(0.042)	(0.055)
Constant	10.991^{***}	10.869^{***}	10.807^{***}	10.734^{***}	10.601^{***}	10.834^{***}
	(0.030)	(0.042)	(0.041)	(0.064)	(0.078)	(0.114)
Observations	619	552	500	449	378	378
Adjusted R^2	0.546	0.470	0.537	0.535	0.509	0.362
Year FE	✓	✓	✓	✓	✓	✓
Other Controls	✓	✓	✓	✓	✓	✓
Job FE						

Notes: Online Appendix Table OA.7 includes estimates from an earnings regression run on the sample of students with identifiable caste names who remain in contention for some job at a given job search stage. For example, column (1) includes the sample of all students with identifiable caste names who apply for jobs recruiting from the college's job fair. Similarly, column (2) includes the sample of all students with identifiable caste names who remain in contention for some job at the written test round, and so on. The dependent variable is log earnings. Each column is a separate regression and includes controls. Note that the earnings regressions do not include job fixed effects because there are no caste differences in pay for identical jobs. Bias-adjusted coefficients are calculated using Oster (2019). Following Oster's recommended benchmarks, I set $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$, implying that unobservables are assumed to have the same degree of selection as observables, and the total explanatory power can increase by up to 30% beyond the observed R^2. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table OA.8: **Caste Differences in Mean Salary at Each Job Search Stage in the Non-Identifiable Sample**

	Dependent variable:					
	log (Salary)					
	App	Test	Debate	Interview	Offer	Choice
	(1)	(2)	(3)	(4)	(5)	(6)
Disadv. Caste (OLS)	0.002	0.0004	−0.003	−0.004	−0.100***	−0.106***
	(0.005)	(0.007)	(0.007)	(0.011)	(0.014)	(0.019)
Disadv. Caste (Oster)	0.004	0.019***	0.013	0.024**	−0.075***	−0.076***
	(0.006)	(0.007)	(0.008)	(0.011)	(0.015)	(0.018)
Constant	11.041***	10.913***	10.857***	10.766***	10.685***	10.869***
	(0.012)	(0.017)	(0.016)	(0.024)	(0.033)	(0.043)
Observations	3,386	3,066	2,766	2,582	2,149	2,149
Adjusted R^2	0.584	0.500	0.552	0.556	0.486	0.421
Year FE	✓	✓	✓	✓	✓	✓
Other Controls	✓	✓	✓	✓	✓	✓
Job FE						

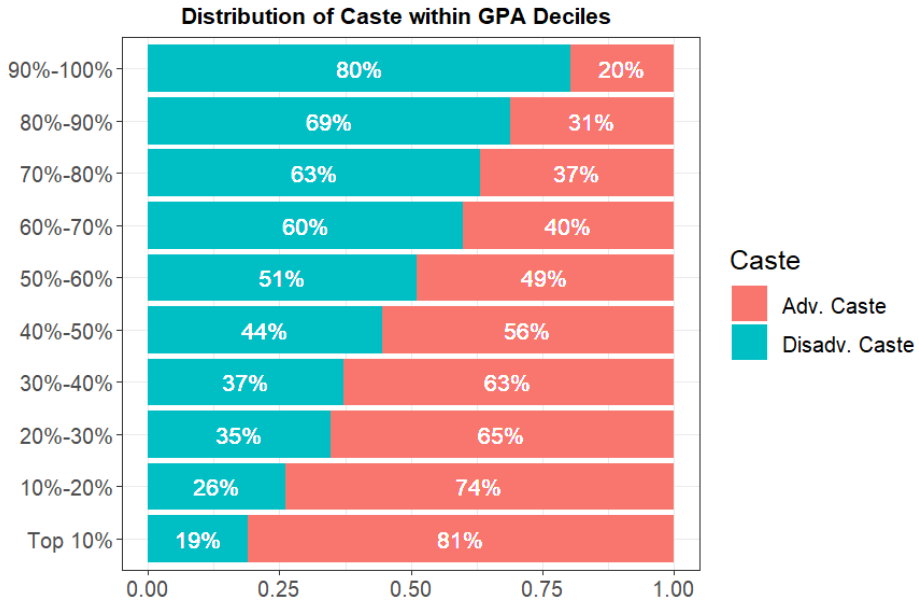
Notes: Online Appendix Table OA.8 includes estimates from an earnings regression run on the sample of students with non-identifiable caste names who remain in contention for some job at a given job search stage. For example, column (1) includes the sample of all students with non-identifiable caste names who apply for jobs recruiting from the college's job fair. Similarly, column (2) includes the sample of all students with non-identifiable caste names who remain in contention for some job at the written test round, and so on. The dependent variable is log earnings. Each column is a separate regression and includes controls. Note that the earnings regressions do not include job fixed effects because there are no caste differences in pay for identical jobs. Bias-adjusted coefficients are calculated using Oster (2019). Following Oster's recommended benchmarks, I set $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$, implying that unobservables are assumed to have the same degree of selection as observables, and the total explanatory power can increase by up to 30% beyond the observed R^2. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table OA.9: **Caste Differences in GPA Between Those Who Left the First Job Versus Those Who Did Not**

	Dependent variable:		
	log (GPA)		
	(1)	(2)	(3)
Disadv. Caste	-0.065^{***}	-0.065^{***}	0.010^{***}
	(0.004)	(0.004)	(0.002)
Left Job	-0.072^{***}	-0.072^{***}	-0.004^{**}
	(0.003)	(0.003)	(0.002)
Left Job * Disadv. Caste	0.047^{***}	0.047^{***}	-0.004
	(0.005)	(0.005)	(0.003)
Constant	2.188^{***}	2.182^{***}	2.198^{***}
	(0.003)	(0.005)	(0.028)
Year FE	✓	✓	✓
Major		✓	✓
Degree		✓	✓
Job FE			✓
Observations	2,489	2,489	2,489
Adjusted R^2	0.260	0.265	0.822
Note:			*p<0.1; **p<0.05; ***p<0.01

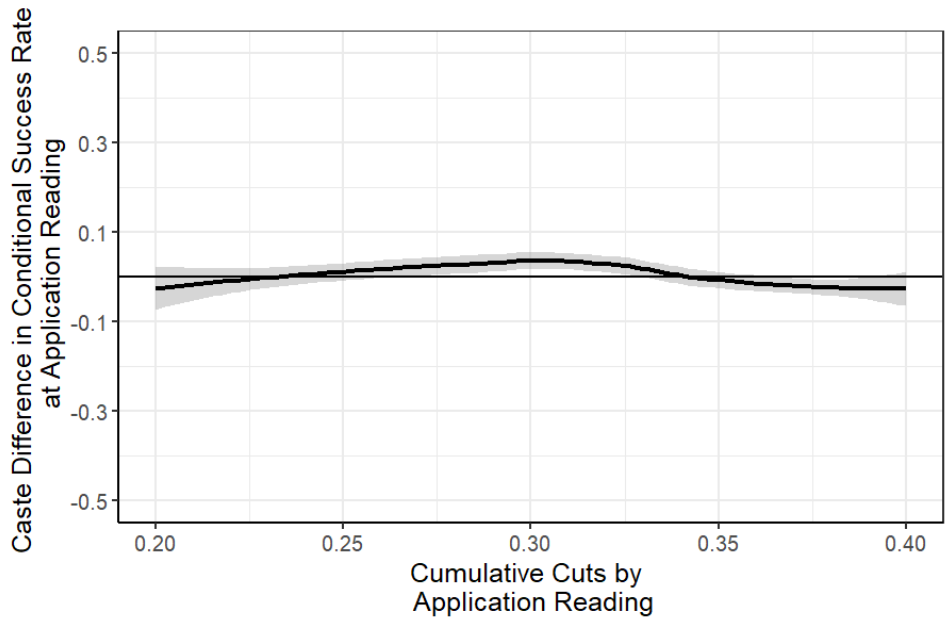
Notes: Online Appendix Table OA.9 shows the caste differences in GPA between those who left the first job versus those who did not. Column (1) only includes year fixed effects. Column (2) adds major and degree as additional controls. Column (3) additionally controls for job fixed effects.

Figure OA.1: **Common Support Within Each College GPA Decile**

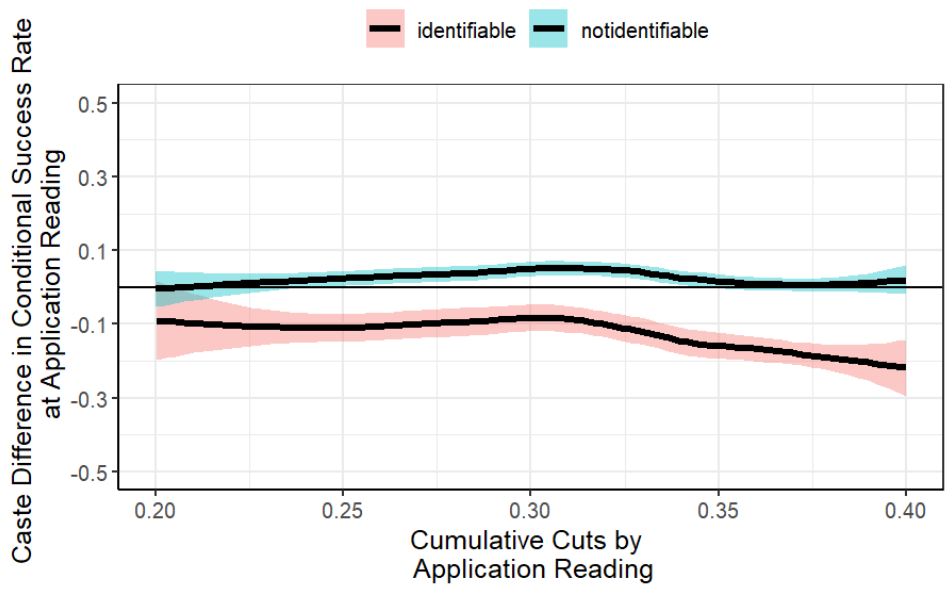


Notes: Online Appendix Figure OA.1 shows that there are students from both disadvantaged and advantaged castes within each college GPA decile. College GPA is normalized within each college degree and year.

Figure OA.2: **Caste Penalty in the Conditional Success Rate At the Application Reading Round in the Identifiable and the Non-Identifiable Sample**



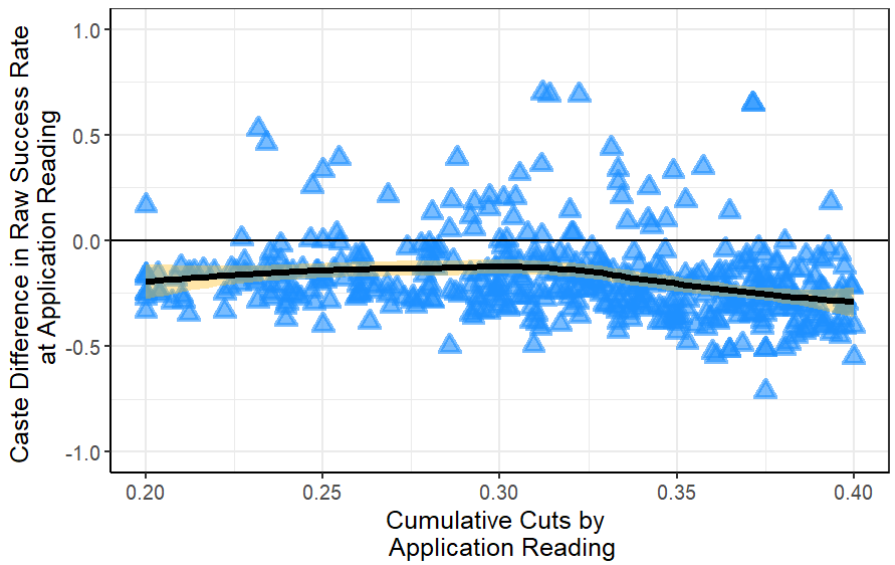
(a) Caste Penalty in Conditional Success Rates versus Cumulative Cuts by Application Reading in the Full Sample



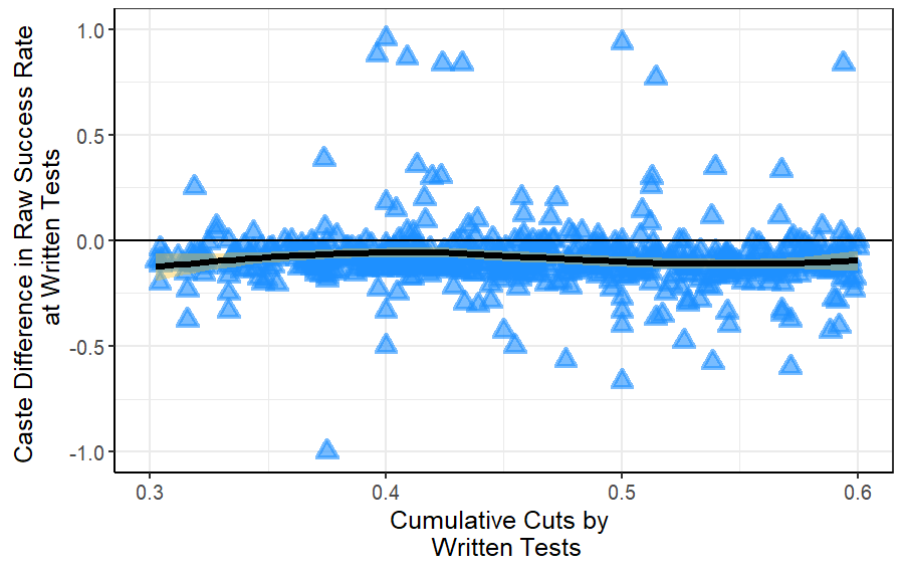
(b) Caste Penalty in Conditional Success Rates versus Cumulative Cuts by Application Reading in the Identifiable and Non-Identifiable Sample

Notes: Figure OA.2 shows the caste difference in the conditional (i.e., with controls) success rate against the cumulative cuts made by jobs by the end of the application reading round. Smoothed conditional means (thick black lines) are calculated using the "loess" function in R. The shaded areas in gray, blue, and red denote 95% confidence intervals. The thin solid horizontal line denotes a y-intercept of 0. Panel (a) is constructed from the full sample. Panel (b) is constructed separately for the sample with identifiable (red) and nonidentifable (blue) caste names. The x-axis shows the cumulative proportion of students cut by jobs by the end of the application reading round. The y-values are calculated by estimating, *for each job*, a standard linear probability model with controls and then plotting the coefficient on disadvantaged caste from these individual regressions on the y-axis. Each regression includes controls for student characteristics such as degree, major, college GPA, entrance exam score, and cohort year.

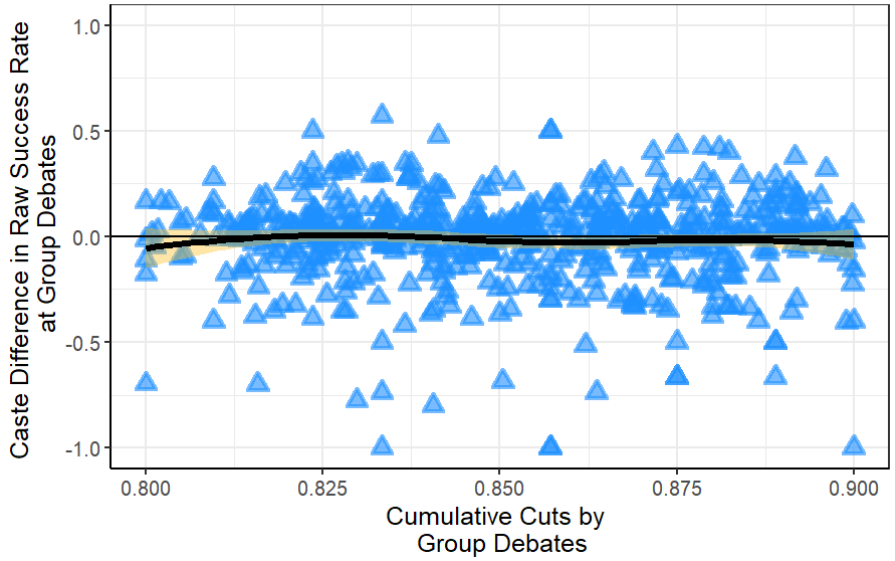
Figure OA.3: **Caste Penalty in the Uncontrolled Success Rate At Each Screening Round**



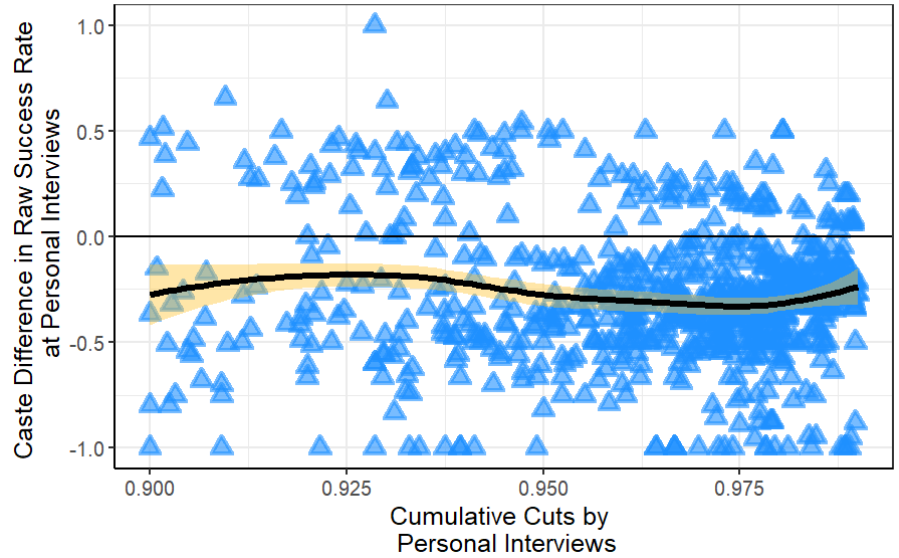
(a) Caste Penalty in Raw Success Rates versus Cumulative Cuts by Application Reading



(b) Caste Penalty in Raw Success Rates versus Cumulative Cuts by Written Tests



(c) Caste Penalty in Raw Success Rates versus Cumulative Cuts by Group Debates



(d) Caste Penalty in Raw Success Rates versus Cumulative Cuts by Personal Interviews

Notes: Figure OA.3 shows the caste difference in the "raw" (i.e., without any controls) success rate against the cumulative cuts made by jobs by the end of each screening round. Each blue triangle denotes a job. Smoothed conditional means (thick black lines) are calculated using the "loess" function in R. The shaded areas in yellow denote 95% confidence intervals. The thin solid horizontal line denotes a y-intercept of 0. The x-axis shows the cumulative proportion of students cut by jobs by the end of each screening round. The y-axis shows the uncontrolled caste penalty in the success rate at each screening round.