Making the Elite: Class Discrimination at Multinationals

Soumitra Shukla*

July 24, 2025

Abstract

What drives socioeconomic disparities in elite job access, and where do these inequalities emerge? Measuring class discrimination is challenging, as status must often be inferred from subtle markers like accent, education, or networks. India's caste system provides an explicit social class measure, enabling precise identification of how status affects elite job access. Analyzing detailed screening data from over 1,000 jobs at primarily U.S. and European multinationals hiring in India, this paper provides the first systematic quantification of class bias in elite recruitment, identifying both its sources and mechanisms. Personal interviews assessing candidate "fit" drive nearly 90% of the caste earnings penalty (0.16 standard deviations). However, disadvantaged caste hires achieve 20% higher promotion rates, suggesting these disparities cannot be justified by ability differences. To formally distinguish discrimination from unobserved ability differences, I develop a theoretical framework utilizing plausibly exogenous variation in screening selectivity across jobs. The framework predicts that, absent discrimination, the promotion advantage of disadvantaged caste hires should shrink in highly selective jobs. Instead, persistently higher and stable promotion rates of disadvantaged hires across all selectivity tiers are most consistent with discrimination. These patterns persist across both client-facing and non-client-facing roles, suggesting bias rather than statistical or customer-driven motives. My findings reveal how subjective "fit" assessments can undermine both diversity and efficient talent allocation in elite hiring.

*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, Peter Arcidiacono, Michael Dinerstein, Arpit Gupta, Patrick Kline, Scott Kominers, Eliza McNay, Suanna Oh, Evan Rose, Kamalini Ramadas, Jung Sakong, Rob Valletta, Jon Williams, seminar participants at the Yale labor workshop, the Chicago FED, Chicago Harris, 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. I used OpenAI's ChatGPT and Anthropic's Claude for light copy-editing.

"People with advantages are loath to believe that they just happen to be people with advantages." —C. Wright Mills, *The Power Elite* (1956)

1 Introduction

Who gains access to top jobs, and how do hiring practices shape inequality in elite recruitment? While extensive research documents labor market discrimination (e.g., Altonji and Blank, 1999), we lack precise evidence on *where*, *how*, and *how much* hiring discrimination occurs in high-stakes, real-world settings like multinationals.

Measuring discrimination is especially challenging when the relevant marker is social class rather than race, gender, or ethnicity. In many settings, including the U.S. and the U.K., status is often inferred from overlapping cues—accents, school pedigrees, inherited networks, income (Weber, 1922; Lamont and Lareau, 1988; Bisin and Verdier, 1998; Lareau, 2003; Durlauf, 1999, 2008; Khan, 2011, 2012; Ridgeway, 2019; Stansbury and Schultz, 2023; Stansbury and Rodriguez, 2024a). Yet these markers often provide ambiguous signals, complicating precise measurement: income can mislead, while accent and educational pedigree can reveal more (Bourdieu, 1986). For example, studies in the U.K. show that a high-earning executive who retains a regional working-class accent is often rated as lower status than a lower-paid professional speaking Received Pronunciation or listing a prestigious boarding school (Savage, 2000; Savage et al., 2013; Savage, 2015; Laurison and Friedman, 2019; Levon et al., 2021).[1] More generally, evaluators choosing among similarly qualified candidates often rely on holistic assessments of who "fits"—referring to candidates' cultural alignment, interpersonal rapport, perceived similarity in values or backgrounds, expectations regarding team dynamics, and other job-specific competencies (Lamont, 1992; Rivera, 2012, 2015).

Holistic evaluations that emphasize "fit" often reinforce class advantages under the guise of comprehensive assessment, encoding deep-seated cultural expectations (Bisin and Verdier, 2011). Test-optional policies reduce access for high-achieving, low-income students (Sacerdote et al., 2025), legacy preferences at selective U.S. universities favor the wealthy despite weak performance links (Chetty et al., 2023), and "personality" interviews used in selection for India's elite civil service penalize lower-status candidates while adding little predictive value (Bhavnani and Lee, 2021). Even outside of these specific examples, decades of evidence show holistic assessments predict no better than standardized metrics, *even among experts*, while systematically disadvantaging lower-status applicants (Meehl, 1954; Ægisdóttir et al., 2006; Kuncel and Hezlett, 2007; Murphy et al., 2009; Rivera, 2012; Bol et al., 2014; Kuncel et al., 2013, 2014;

[1] Received Pronunciation, sometimes called "BBC English" or "Queen's English," is the historically prestigious accent of English traditionally associated with the British upper classes and elite institutions (Savage et al., 2013).

Laurison and Friedman, 2019; Stansbury and Rodriguez, 2024a). Thus, the very mechanism meant to refine selection often obscures merit (Comerford et al., 2022).

Establishing that fit assessments lead to social class discrimination requires precise measurement of class itself—a challenge that motivates my focus on the Indian caste system, which provides an unusually comprehensive marker of social stratification. While caste is uniquely Indian, it embodies the broader concept of social class that Bourdieu (1986) conceptualizes as a joint stock of four capital forms: economic (money, property rights), social (durable networks), cultural (embodied traits like accent, etiquette, and social ease; leisure activities, familiarity with cultural references and shared experiences; institutionalized credentials from educational pedigree and family background), and symbolic (prestige from public recognition of other capitals). Unlike the more fluid manifestations of class in other contexts, India's caste system crystallizes these capital forms into a relatively stable, administratively codified hierarchy (Beteille, 1965, 1969). This codification provides rare precision for measuring class discrimination in elite jobs, which disadvantaged castes remain largely excluded from despite comprising 70% of India's billion-plus population (Bayly, 2008).

Leveraging this measurement advantage, I examine recruitment for India operations at primarily U.S. and European multinationals to make three contributions: First, I quantify social-status discrimination in elite recruitment. Second, I identify personal interviews, which assess candidate fit, as the primary mechanism revealing and penalizing social status. Third, I show that class bias—not statistical or customer discrimination—drives these disparities.[2] To the best of my knowledge, each of these contributions is novel in the literature.

Personal interviews penalize social status, driving nearly all caste penalties in hiring. In my sample of over 1,000 jobs, each position involves four screening rounds progressing from standardized to subjective assessments: application reading, aptitude tests, group debates that test socio-emotional skills, and, finally, personal interviews assessing candidate "fit." These jobs all recruit from an elite Indian college, with nearly half the students belonging to disadvantaged castes due to *mandated affirmative action* (Madheswaran, 2008). The unconditional caste earnings gap is 23%, narrowing to 10% after controlling for pre-college skills, academic performance, and other employer-relevant factors. Since salaries are explicitly *identical across castes* for identical roles and *non-negotiable*, one can further decompose these gaps. This is done by counterfactually assigning posted salaries to the set of jobs each caste group is in contention for at each screening stage before offers are accepted. Personal interviews alone account for nearly 90% of the conditional earnings penalty (approximately 9% of average earnings, or 0.16 standard deviations) and about 95% of the nearly 30 percentage point caste penalty in hiring success rates, highlighting their role in penalizing social status.

[2]Such discrimination violates Article 15 of the Indian Constitution, limiting mobility and misallocating talent.

Personal interviews reveal social status. While some surnames indicate caste, most do not convey caste signals in urban-educated settings (e.g., elite colleges typically located in urban centers) due to migration and geographical variation in naming conventions (Mamidi, 2011; Jodhka, 2015; Wiltshire, 2020). Among the 15% of students in my sample with distinctly caste-identifiable surnames, penalties emerge immediately at the first screening round (application reading). For the remaining 85% with caste-ambiguous surnames, no early penalty appears—disparities emerge only after personal interviews. This differential timing strongly suggests that the penalty occurs precisely when caste becomes visible to employers, with personal interviews serving as the primary revelation mechanism for social status.

Interviews may provide information, but cannot justify observed disparities. Disadvantaged caste hires consistently outperform advantaged peers by around 20% in promotion rates, even in highly selective roles, strongly indicating disparities stem from discrimination rather than unmeasured skill deficits. To formalize this interpretation, I develop a theoretical framework that leverages plausibly exogenous variation in screening selectivity (the fraction of the applicant pool hired) across jobs. The framework yields two key implications: first, if disadvantaged candidates genuinely lacked skills, highly selective pre-interview screening stages would show a caste penalty, which they do not; second, if interviews accurately filtered weaker candidates, disadvantaged hires would not consistently outperform advantaged peers across selectivity tiers. As selectivity increases, the average hired candidate's quality approaches marginal quality (hiring cutoff). If jobs were using the same hiring cutoffs regardless of caste, promotion differences between castes should shrink as job selectivity increases. Thus, persistently higher and stable promotion rates among disadvantaged hires across all selectivity levels suggest discrimination rather than productivity differences.

Analysis of distinctive versus non-distinctive surnames, combined with the promotion edge of hired disadvantaged castes across job selectivity tiers, strongly indicates that hiring disparities reflect *disparate treatment* (explicit caste-based discrimination). Moreover, disadvantaged caste hires retain their promotion advantage across both client-facing and non-client-facing jobs, suggesting that class bias (either taste-based discrimination or biased beliefs) rather than statistical or customer-driven discrimination drives the caste hiring penalty.

This paper contributes to labor, personnel, and organizational economics by quantifying social-status discrimination in elite hiring and identifying its mechanisms. This paper advances the still-nascent literature on class discrimination, which lags behind research on race, gender, and ethnicity (Durlauf, 2006; Rivera, 2012, 2015). The paper is most closely related to Stansbury and Schultz (2023), who show that access to economics PhD programs is disproportionately skewed toward those from higher socioeconomic backgrounds, and Stansbury and Rodriguez (2024a), who document sizable class gaps in career progression within U.S. academia; Rivera (2012), who argues that "cultural matching" in elite-firm interviews could engender class

disparities; Bhavnani and Lee (2021), who show that disadvantaged caste civil service officers score markedly lower in "personality" interviews, yet still outperform advantaged caste recruits; and Zimmerman (2019), Michelman et al. (2022), and Barrios-Fernández et al. (2024), who demonstrate the role of elite education in reproducing class advantage through what Bourdieu (1986) conceptualizes as peer-driven social and cultural capital. No prior work, however, simultaneously quantifies the magnitude of class discrimination, locates it in the hiring pipeline, and isolates its mechanisms.

As traditional social-status markers become less visible across societies, class discrimination often operates through *social interactions*—such as interviews—where subtle cues can reveal and penalize social background (Lenski, 1954; Grusky, 2005; Savage et al., 2015; Rivera, 2015; Clark, 2014; Rose, 2023). Indeed, what Wilkerson (2020, pp. 22-23) calls America's "unseen skeleton, a caste system that is [central] to its operation" functions as a "wordless usher in a darkened theater [...] guiding us to our assigned seats"—aptly capturing how class hierarchies persist through subtle social choreography. This study demonstrates how subjective fit evaluations can preserve discrimination even within ostensibly meritocratic systems, often enabling biases while maintaining plausible deniability (Castilla and Benard, 2010; Reskin, 2012; Castilla, 2015; Kang et al., 2016; Castilla and Ranganathan, 2020). Such disparities manifest in the subjective interpretation of performance signals across various domains, such as gender disparities in academic co-authoring, experimental group projects, and medical referrals (Sarsons, 2017; Sarsons et al., 2021; Sarsons, 2023). These patterns likely extend beyond visible characteristics to concealable stigmatized identities: LGBTQ+ status and other intersectional attributes (Chow and Knowles, 2016; Eames, 2024).

This paper traces how personal interviews—a widely adopted practice at multinationals—shape labor market inequality by affording hiring managers the opportunity to discriminate on social status, a topic that has received limited attention in economics. While economic studies extensively document discrimination at early stages like résumé screens (e.g., Kline et al., 2022), few examine bias within interview rooms where subjective evaluations occur. Meta-analyses suggest additional discrimination often emerges during face-to-face interactions, as race and gender gaps persist or widen from initial callbacks to final offers (Quillian et al., 2020). By pinpointing when class penalties crystallize and linking them to post-hire performance, this study provides novel evidence on a ubiquitous yet understudied hiring practice while highlighting limitations of interventions like résumé anonymization that may merely shift bias to later stages.

Another central contribution illuminates managers' critical role in perpetuating class discrimination through hiring decisions, revealing both cultural biases and structural agency problems. Recent literature documents substantial managerial influence on firm performance and culture (Turco, 2010; Lazear et al., 2015; Sadun et al., 2017; Hoffman and Tadelis, 2021; Es-

pinosa and Stanton, 2024; Fenizia, 2022; Sadun, 2023; Benson and Shaw, 2024; Minni et al., 2025), particularly through assigning workers to teams and training investments (Minni, 2024b; Sadun, 2023; Sadun et al., 2024). This paper uniquely integrates Bourdieu's concept of social class—especially its social and cultural capital dimensions—to illuminate biases in managerial assessments of who "fits" (Bourdieu, 1986). Additionally, this paper highlights agency problems in hiring: the separation of hiring authority from performance accountability and managers' reliance on overconfident private judgments (Hoffman et al., 2018) or productivity-irrelevant cues (Benson et al., 2023). Hiring managers in my setting rarely supervise recruits directly, creating incentives to prioritize perceived "fit" over long-term potential. Such subjective, short-term decisions systematically produce sub-optimal organizational outcomes, as documented by extensive research (Gibbons and Roberts, 2014; Sadun et al., 2017; Benson and Shaw, 2024; Hoffman and Stanton, 2024).

Finally, this paper reveals how subjective screening practices that provide room for status-based discrimination can erode affirmative action gains, undermining the socioeconomic mobility that elite education is designed to create. Elite colleges like the studied Indian institution, with a talent pool comparable to that at many Ivy League universities (Martellini et al., 2024) and affirmative action policies in admissions (Madheswaran, 2008), serve as gateways to prestigious careers (Jodhka and Naudet, 2019). Although similar affirmative action policies worldwide substantially increase access to elite colleges for students from lower social backgrounds, research indicates elite education disproportionately benefits privileged students, perpetuating class advantage across generations (Zimmerman, 2019; Barrios-Fernández et al., 2024). My findings highlight the need to integrate affirmative action in college admissions with hiring reforms downstream, emphasizing comprehensive approaches for socioeconomic mobility—particularly as social class co-evolves with cultural and institutional factors that shape economic outcomes while simultaneously maintaining class boundaries (Bisin et al., 2023; Bisin and Verdier, 2025).

2 Institutional Setting and Key Definitions

This section provides essential background on India's caste system as a measure of social status, defines key terms used throughout the analysis, describes the elite college job fair studied in this paper, and explains why this setting is representative of elite hiring in India.

2.1 Caste as Social Status Under Affirmative Action

India's affirmative-action policy framework treats caste as a unified marker of social status, bundling economic, social, and educational advantages into two administratively codified categories: "advantaged" (upper) and "disadvantaged" (lower). These policies apply only to public

universities and government jobs, not to private multinationals. For historical background on caste classification and affirmative action policy, see Online Appendix Section A.

2.2 Key Terms: Disadvantaged Castes, Elite Jobs, and Elite Colleges

Below, I define four key terms used throughout the analysis:

a. **Disadvantaged castes.** Following the Indian government's official classifications, the term "disadvantaged castes" encompasses individuals belonging to Scheduled Castes (SCs), Scheduled Tribes (STs), or Other Backward Classes (OBCs).

b. **Job.** A "job" refers specifically to a job designation within a firm. For example, if Google recruits both a product manager and a software engineer, these constitute two distinct jobs. Section 3.4 explains why a job typically corresponds to a firm-designation pair, rather than a particular office location.

c. **Elite jobs.** An "elite entry-level job" refers to a formal entry-level position at multinationals, major national corporations, or unicorn startups (privately-held startups valued at over $1 billion). Throughout this paper, the terms "elite entry-level jobs" and "elite jobs" are used interchangeably.

d. **Elite colleges.** "Elite colleges" include institutions consistently ranked among the top 30 in their respective fields—such as science, engineering, arts and humanities, social sciences, and business—by *India Today*, an Indian equivalent of the *U.S. News & World Report*. Several elite Indian colleges place among the top colleges worldwide, according to a cross-country analysis of college quality in Martellini et al. (2024), underscoring their international standing.

2.3 Structure of the On-Campus Job Fair

The elite college's job fair occurs *entirely on campus* and comprises five phases: (1) pre-placement, (2) return offers from internships, (3) student and employer registration, (4) placement, and (5) post-placement ("aftermarket").

a. **Pre-Placement phase.** Firms are invited by the placement office prior to June. Between June and mid-August, recruiters deliver pre-placement talks, gauge student interest, and finalize job roles.

b. **Return offers from internships.** By late August, firms must extend pre-placement offers (PPOs) to students who interned with them during the summer, following deadlines set by the placement office. The placement office withdraws students who accept PPOs from the main placement process. *Unlike in other contexts such as the U.S., these internship return offers play a minor role*: fewer than 5% of students accept PPOs and exit the placement process, leaving over 95% to actively participate in the main job fair.

c. **Student and employer registration.** Students register for placement by late August, while employers submit detailed registration forms by early September, specifying job roles, compensation, and vacancies. Compensation and benefits are *non-negotiable* during the course of the on-campus job fair, and are strictly enforced by the placement office through verification of job offer letters. Non-compliance can result in restrictions on future campus visits by firms.[3]

d. **Placement phase.** Starting in mid-September, students apply for jobs. Employers initially screen applications, followed by written aptitude tests and group debates assessing socio-emotional skills from September to early December. Final-round personal interviews occur between December and January, after which about two-thirds of the total student cohort secure full-time jobs, including those who accepted summer internship return offers.

e. **Post-Placement phase ("aftermarket").** From late January until graduation in July, the placement office operates an "aftermarket" for students without initial job offers, comprising roughly one-third of the cohort. These unplaced students typically pursue entrepreneurship, graduate studies, competitive exam preparation (such as civil services or MBAs), or find alternative employment. Aftermarket employers—typically lower-paying or recent startups—use less standardized recruitment processes than firms participating in the main placement cycle.

2.4 Why the Studied Job Fair is Representative

The studied job fair is representative of how elite Indian graduates transition into elite jobs. Nearly all top-ranked Indian colleges use similar standardized placement procedures, graduates from these institutions predominantly pursue elite private-sector positions, elite employers recruit almost exclusively from the same set of campuses, and student participation rates in campus job fairs are near-universal. For evidence supporting these claims, see Online Appendix Section B.

[3]Such rules and timelines are common across elite Indian colleges and are available on their websites. For instance, see Central Placement Cell (Delhi University).

3 Data Overview

This section uses a dataset of roughly 4,000 graduates in technical majors—about half from disadvantaged castes—who competed for more than 1,000 job positions requiring strong technical backgrounds (e.g., software engineering, data analytics, quantitative consulting) through a standardized four-stage screening pipeline. The technical demands of such roles coupled with structured evaluation should leave little room for late-stage subjectivity and social class considerations, making any remaining class-based advantages especially informative. We describe student and job characteristics, recruitment procedures, career outcomes (promotions and tenure), and key data limitations.

3.1 Institutional Source and Completeness of Recruitment Data

Data are obtained directly from the placement office of an elite Indian university (anonymized for confidentiality). The university's standardized job fair serves as a benchmark for placement processes at comparable institutions nationwide (e.g, Jodhka and Naudet, 2019).

3.1.1 Centralized Job Portal

Students register through a centralized online job portal managed by the placement office, uploading résumés and directly entering standardized information including college GPA, college entrance exam scores, coursework, and internship experiences. Firms submit detailed profiles specifying compensation, job descriptions, and qualifications (Section 2.3). Both student and job profiles are mutually visible via the portal. Throughout recruitment, employers directly upload candidate shortlists from each recruitment stage—applications, intermediate screening tests, interviews, and final job offers onto the portal, ensuring systematic documentation and data consistency.

3.1.2 Standardized Recruitment Rules and Mandatory Participation for Registered Students

Recruitment at elite Indian colleges is tightly standardized: students who enter the campus placement fair are *mandated* by the placement office to participate in all four screening rounds—application screening, written aptitude tests, group debates, and personal interviews (Section 3.6). Offer-letter salaries and benefits are likewise verified against the non-negotiable compensation schedules filed by firms (Section 2.3). Such standardized rules enhance transparency and empirical rigor.

3.2 Student Sample: Time Span, Degrees, Demographics, and Caste Representation

The dataset spans four placement cycles at an elite Indian college (exact years anonymized for confidentiality). Table 1 shows that there are 4,164 students pursuing four degrees: Bachelor of Technology (B.Tech.), Dual Degree (integrating undergraduate and postgraduate coursework), Master of Technology (M.Tech.), and Master of Science (M.S.).

Admissions based solely on standardized national entrance exams. Admissions are *exclusively* based on standardized national-level exams with caste- and major-specific cutoffs.

Near-equal caste representation via affirmative action. Due to caste-based affirmative action in admissions (Section 2), disadvantaged and advantaged castes are represented at near-parity *within each major*, and therefore within most degrees.[4]

Caste self-reporting. Students self-report caste through official administrative documents, enabling accurate tracking by both the college and researcher.

Predominantly male sample. Approximately 90% of students are male, typical of elite technical colleges in India (Datta, 2017).

[4]In some degrees, particularly the Master of Science (M.S.), disadvantaged-caste seats often remain partially unfilled due to eligibility constraints or fewer qualified applicants, consistent with patterns documented in other research (e.g., Datta, 2017).

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 student counts by caste and college degree. B.Tech. stands for Bachelor of Technology, M.Tech. for Master of Technology, and M.S. for Master of Science. A Dual degree integrates undergraduate and post-graduate studies and is completed one year after the four-year B.Tech.

3.3 Student Characteristics and Caste Differences

Student characteristics include 10th and 12th standard national exit exam scores, college entrance exam scores, college GPA, and prior labor market experience. Two key patterns emerge:

1. **Significant caste differences in academic performance.** Large caste disparities exist in college GPA and college entrance exam scores (Table 2), although students from both caste groups span all deciles of such metrics (e.g., common support for college GPA is shown in Online Appendix Figure OA.1).

2. **Similar internship patterns across castes.** The dataset includes detailed internship information (e.g., duration, pay, employment sector). Caste is only weakly correlated with internship sector or experience. For instance, among Bachelor's and Dual Degree students, internships across castes show comparable sector distributions (25% technology, 37% consulting, and 38% manufacturing). Similar results hold for postgraduate degrees (not shown). Such patterns likely reflect that return internship offers are rare (Section 2.3), making internship choices primarily driven by student sector preferences rather than caste-related pay or opportunities.

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 reports pre-college skills (college entrance exam scores, 10th, and 12th standard national exit exam scores) and college GPA. B.Tech. stands for Bachelor of Technology, M.Tech. for Master of Technology, and M.S. for Master of Science. A Dual Degree integrates undergraduate and post-graduate studies and is completed one year after the four-year B.Tech. All scores are normalized to zero mean and unit standard deviation within each degree and year. Entrance exam scores (originally expressed as ranks) are renormalized so higher numbers are better. Caste differences shown in standard deviation units. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$. Adv. Caste and Disadv. Caste denote advantaged and disadvantaged castes.

3.4 Characteristics of Multinational Private-Sector Jobs

This subsection summarizes key job characteristics, including sector distribution, salaries, recruitment procedures, and managerial involvement.

Sector and employer overview. The dataset includes over 1,000 private-sector jobs, mostly offered by multinationals headquartered in the U.S. or Europe (Table 3). Around 75% of roles are non-client-facing (e.g., software engineering) and 25% client-facing (e.g., quantitative consulting), with technically-oriented, non-client-facing roles commanding premium salaries at this elite technical college.

Excluding public-sector jobs. Public-sector roles (about 2%; Table 3) are excluded due to distinct salary structures and recruitment procedures (Seventh Central Pay Commission, 2016).

Standardized, non-negotiable salaries and benefits. For a given college degree, salaries and benefits are standardized, non-negotiable, and uniform for each job role, independent of individual characteristics like caste, gender or major (or any other candidate-specific factors)—a feature that strengthens identification by ensuring any earnings gaps reflect differential *job access* rather than idiosyncrasies of wage setting. Campus recruiters usually have no discretion to vary pay across locations or candidates, as firms adopt nationwide salary standards, recruiting identically across elite campuses, consistent with how multinationals centralize entry-level compensation and recruitment procedures at their headquarters, with local branch offices implementing these centralized directives (e.g., Bloom et al., 2012; Minni, 2024a,b; Sadun et al., 2024).

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 reports the distribution of jobs between public and private sectors, domestic and foreign-headquartered (primarily US or Europe) firms, and client-facing versus non-client-facing roles. Average salaries (USD PPP) are also shown for each group.

Job location. About 99% of private-sector jobs are India-based. Non-Indian positions use different recruitment teams and compensation packages—for example, (Google, software engineer, Palo Alto) differs from (Google, software engineer, Mumbai) in pay as well as the team that conducts recruitment. In contrast, Indian positions are recruited jointly across all Indian offices with standardized compensation, making (Google, software engineer, Mumbai) and (Google, software engineer, Bangalore) effectively identical. Office assignment for Indian positions occurs later, near joining date, depending upon location-specific needs of the employer. Thus, a "job" is defined as a firm-designation pair for Indian positions but includes location for non-Indian positions (see also Section 2.2).

Recruitment conducted by Indian hiring managers separately from direct supervisors. Because nearly all private-sector jobs are located in India, recruitment is conducted by specialized Indian hiring managers. These managers typically remain distinct from eventual direct supervisors due to logistical constraints posed by placement processes of elite colleges that must be entirely conducted *on-campus*, including intermediate screening tests and final interviews that are personally supervised or conducted by recruiters in college-assigned rooms.

Internship return offer policy. Students accepting internship return offers are withdrawn from the main placement cycle by the placement office and are excluded from analysis (Section 2.3), with implications of this exclusion discussed in Section 4.4.

3.5 Quantifying Job Selectivity at Each Recruitment Stage

This subsection summarizes the selectivity of jobs using two key concepts—*conditional cuts*, the fraction of applicants eliminated *at* a given screening round, and *cumulative cuts*, the total fraction eliminated *by* a given screening round. Table 4 reports these metrics. For instance, a job eliminating 50% of applicants each round has a 50% conditional cut at each stage, but a cumulative cut of 75% after two rounds. Typically, successful candidates rank within the top 5% of initial applicants. Occasionally, highly selective roles extend no offers.

Table 4: **Variation in Conditional and Cumulative Cuts 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 conditional and cumulative cuts across jobs. Conditional cuts show the fraction of candidates eliminated at each stage; cumulative cuts show the total fraction eliminated by that stage.

3.6 Four Sequential Recruitment Stages: From Objective to Subjective Evaluations

For *all jobs* in the analytic sample, recruitment progresses uniformly from relatively objective assessments to subjective evaluations through four sequential stages, with candidates eliminated at each round (Table 4):

1. **Application reading.** Employers use standardized criteria and GPA cutoffs to initially filter résumés, typically involving limited subjective assessment.

2. **Written aptitude tests.** Standardized online tests objectively evaluate technical, analytical, and critical thinking skills.

3. **Group debates.** Structured debates, conducted in multiple batches, involve 20–30 students divided into two teams, assessing communication, confidence, and teamwork under standardized conditions, including uniform prompts and roughly equal speaking time per candidate.

4. **Personal interviews.** Also sometimes referred to as "HR interviews," these final-round interviews last about 30 minutes and emphasize subjective evaluations of interpersonal skills, organizational fit, cultural alignment, and role-specific competencies.

Throughout each stage, employers maintain access to candidates' standardized profiles via the centralized portal (Section 3.1.1). Additional details for screening rounds and candidate selection, including the timeline, are discussed in Online Appendix Section C.

Institutional details of the application and screening process.

a) Low-cost applications via centralized portal. A single click on the centralized portal (Section 3.1.1) forwards the same standardized résumé to any additional job, making the marginal cost of applying essentially zero; application behavior is therefore virtually identical across castes (Figure 1).

b) Eligibility criteria restrict applications. Applications are constrained by major, degree, and GPA requirements. For example, software engineering roles often require Computer Science majors, whereas analyst roles typically allow all majors. Eligibility constraints (major, degree, GPA) naturally limit the number of applications per student.

c) Common screening rounds across multiple positions. Employers offering multiple distinct roles frequently administer common screening simultaneously for all applicants (e.g., Nvidia uses a single exam for all advertised profiles).

3.7 Candidate Screening: Summary Statistics

Table 5 summarizes candidate progression through each recruitment stage. On average, jobs receive roughly 209 applications, with approximately 125 advancing to written tests, 102 reaching group debates, 26 participating in personal interviews, and about 6 candidates receiving job offers. This substantial reduction reflects high selectivity at each stage (Table 4).

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 reports mean and median counts per job for applications, written tests, group debates, interviews, and offers. Values calculated by averaging student counts across years per job, then averaging across jobs.

3.8 Data on Job Tenure and Promotions

Career outcomes are measured using tenure and promotions, obtained for nearly all (over 98%) placed students using alumni records[5], also maintained by the placement office and largely curated from an employment-focused online platform plus a few other sources (see Online Appendix Section D for coverage benchmarks).

Classifying promotions. Promotions are defined as movements to higher job titles within the same firm, validated by salary increases and by comparing against typical tenure-to-title progressions using external industry-standard benchmarks (such as Glassdoor, Levels.fyi) that help distinguish them from lateral moves. This approach of classifying promotions is standard in the literature (e.g., Wilmers et al., 2025).

3.9 Data Limitations and Mitigation

Two key limitations include the absence of detailed résumés and screening round performance scores. However, standardized portal data (Section 3.1.1) comprehensively capture many rel-

[5]The missing 2% show no systematic caste differences in observable characteristics (not shown).

evant student characteristics (e.g., GPA, entrance exam scores, internships), ensuring robust empirical analysis. Strategies addressing the second limitation are discussed in Section 4.7.

4 Penalizing Social Status Through "Fit" Assessments

This section presents six core findings demonstrating how personal interviews penalize social status in entry-level hiring at multinationals. First, roughly 90% of caste-based disparities in earnings and success rates arise specifically during subjective, "fit"-oriented personal interviews. Second, these disparities are likely conservative. Third, personal interviews are the primary mechanism through which social status is revealed and penalized. Fourth, caste penalties consistently emerge at personal interviews across nearly all job positions. Fifth, disadvantaged caste hires exhibit systematically higher promotion rates even in highly selective roles, strongly suggesting hiring disparities result from discrimination rather than unobserved skill differences. Finally, these caste penalties align most closely with *bias* rather than statistical or customer-driven discrimination.

4.1 Main Sample and Reliability of Initial Job Placements

As discussed in Section 2.3, approximately two-thirds of students secure full-time positions. Specifically, about 4% accept pre-placement return offers from their summer internships, while about 61% obtain employment via the standardized four-stage on-campus process. The remaining one-third participate in the subsequent *aftermarket* hiring phase.[6]

The analysis focuses on students who secure jobs through the standardized four-stage process, excluding those who accept internship return offers and the roughly one-third who remain unplaced after the main cycle. Many of these unplaced students participate in the "aftermarket" (Section 2.3), where employers—typically smaller startups offering lower salaries—bypass the four-stage pipeline, using less standardized recruitment processes (e.g., only résumé screening and interviews). These exclusions are necessary given the focus on standardized hiring practices across jobs; sample selection implications are discussed in Section 4.4.

All placed students begin employment in the roles offered. These placements are further verified through post-graduation surveys conducted by the placement office and cross-validated against external employment platforms, affirming the reliability of placement records as accurate indicators of initial job outcomes.

[6]Of the total 4,164 students, around 159 (4%) accept internship return offers, leaving 4,005 actively participating. Approximately 2,527 students (61% of the total cohort) secure jobs via the four-stage process, totaling 2,686 students (159 internship returns + 2,527 from placement pipeline), or about 65% placed through the primary placement process. Consequently, roughly 1,478 students (about one-third) enter the aftermarket. See Online Appendix Tables OA.1 and OA.5 for some of these numbers.

4.2 Caste Penalties in Earnings and Hiring Success Rates Emerge at Personal Interviews

Caste disparities in both initial earnings and hiring success rates predominantly arise after personal interviews and are robust to alternative specifications (Section 4.3).

Earnings Gap Decomposition. Disadvantaged castes earn significantly lower initial salaries overall: the unconditional earnings gap is approximately 23%, narrowing to roughly 10% after controlling for pre-college skills, college GPA, prior labor market experience, and other factors. This residual 10% gap is about $\$6,000$ (PPP) or approximately 0.18 standard deviations (see Table 3). Online Appendix Table OA.1 shows the raw and controlled earnings regressions, as well as separate results for client- and non-client-facing jobs.

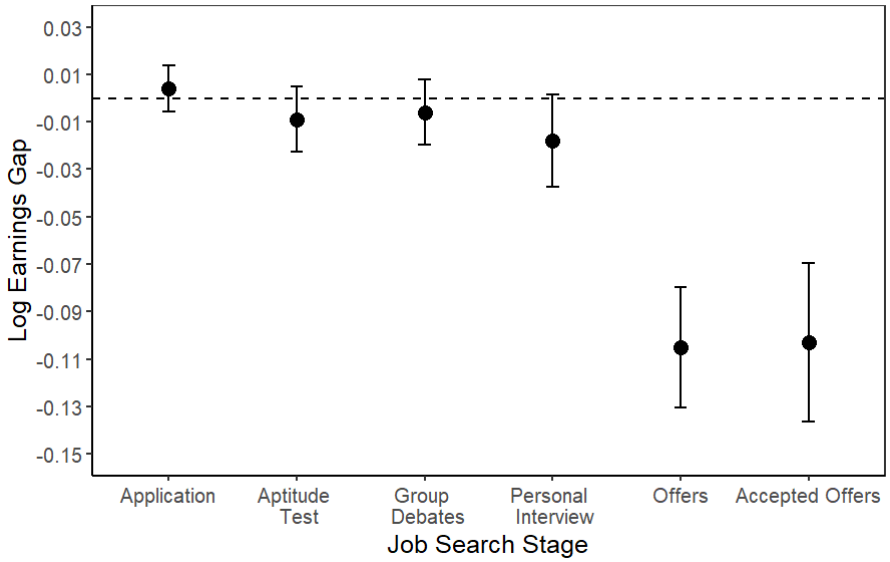
Because jobs have *fixed, non-negotiable* salaries, caste earnings penalties across screening stages reflect differences in *job access* rather than candidate-specific salary variations (Section 3.4). This feature enables clear identification of *where* in the screening process the caste earnings penalty first appears by allowing one to further decompose the earnings gap. This is done by counterfactually assigning posted salaries to the set of jobs each caste group is in contention for at each screening stage before final offers are accepted. The approach quantifies *how much* of the caste earnings penalty accumulates at each round of screening. Formally, I estimate the following regression specification:

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

where $i \in \mathcal{I}$ indexes students, and "Search Stage" $\in$ {Application, Written Tests, Group Debates, Personal Interviews, Offers, Accepted Offers}. Controls include college GPA, college entrance exam scores, prior internships, major, degree, and other characteristics; γ_t denotes cohort-year fixed effects. For each student, the "average job salary" at a given stage is the mean salary across all jobs for which the student remains in contention.

The coefficient β, illustrated in Figure 1, quantifies the caste earnings gap at each recruitment stage. The application reading round shows no statistically significant earnings gap, confirming that low-cost, broad applications (Section 3.6) lead to similar application behavior across castes. The written test and group debate rounds likewise reveal no meaningful gaps. Differential caste preferences over offered jobs also contribute negligibly to the total penalty. The earnings penalty emerges sharply at the personal interview stage, accounting for roughly 90% of the total earnings gap (0.16 standard deviations). Full regression results, including robustness specifications, are presented in Online Appendix Table OA.2.

Figure 1: **The Caste Earnings Penalty**



Notes: Figure 1 shows the log earnings gap across castes at each job search stage. Black dots represent coefficients showing the percentage difference in average salary between disadvantaged and advantaged castes, with 95% confidence intervals. Since salaries are identical and non-negotiable for identical jobs, the earnings gap reflects differences in job access at each stage. All regressions include controls.

Success Rate Gaps by Screening Stage. To provide complementary evidence to the earnings analysis, I examine success rates—the probability of progressing from one recruitment stage to the next *within* each job. Because salaries are fixed, a negligible earnings gap at stage t+1 can occur only if both caste groups have similar odds of advancing past stage t; otherwise they would be in contention for different sets of jobs at stage t+1, resulting in immediate pay differences. Success-rate regressions therefore furnish a direct consistency check on the earnings decomposition in Figure 1 while also accounting for heterogeneity across jobs.

Caste differences in success rates are negligible at the application reading, written test, and group debate rounds (Table 6). The gap opens only after personal interviews, where disadvantaged castes are 30 percentage points less likely to advance, a drop that explains more than 95% of the overall success-rate disparity, aligning with the interview-driven earnings penalty.[7] The same pattern holds among client-facing and non-client-facing jobs (Online Appendix Tables OA.3 and OA.4) and is robust to adjustments for omitted variable bias (Oster (2019)

[7]Based on the estimates in Table 6, disadvantaged caste candidates in our sample are about 70% less likely than their advantaged peers to clear the final interview stage. The risk-ratio—defined as the probability of success for disadvantaged castes divided by the probability of success for advantaged castes—is approximately 0.28, indicating disadvantaged candidates have only 28% of the success rate of advantaged candidates. Thorat and Attewell (2007) report an average 50% shortfall in initial callbacks for disadvantaged groups (risk-ratio $\approx$ 0.52). Although the point estimates differ, both imply that disadvantaged groups face roughly a one-half to two-thirds reduction in success probability, reflecting consistently sized disparities across different hiring contexts.

bounds in Online Appendix Section E).

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)
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 reports caste differences in success rates at each job search stage. Each observation is a student-job-year combination for candidates still in contention at that stage. All regressions include job fixed effects and controls. When firms conduct common screening rounds for multiple positions, each position is treated as a separate job. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

4.3 Robustness and Job-Level Heterogeneity in the Caste Penalty

These findings withstand extensive robustness checks, including a fine-grained *job-level* analysis where I estimate the caste penalty separately for each employer–position pair. Because firms likely vary in written-test content, group debate format, interviewing style, and caste composition, documenting a roughly 30 percentage point caste penalty *within nearly every job* rules out concerns that firm-specific screening methods or other forms of organizational heterogeneity explain the gap. Section 4.7 details these estimates and presents related job-performance evidence based on promotion outcomes.

4.4 Conservative Estimates Due to Sample Exclusions

The main analysis excludes students who do not secure jobs through the full-time placement process (Section 4.1). This group comprises (i) students who enter the post-placement "aftermarket" and (ii) those who accept pre-placement internship return offers, who are then withdrawn by the placement office from the formal placement process (Section 2.3). Disadvantaged caste students among the excluded generally exhibit substantially *lower* average GPAs than comparable advantaged castes (Online Appendix Table OA.5).[8] Because a higher GPA

[8] *Within* any offered job, disadvantaged hires actually have slightly higher GPAs than advantaged counterparts, consistent with the caste hiring penalty (Online Appendix Table OA.6).

typically correlates with higher-paying and more selective job placements, excluding a subset of lower-GPA disadvantaged students from the main analysis understates the true caste disparities.

4.5 Information Observed by Employers at Each Screening Round

Employers continuously access candidates' standardized academic profiles, including GPA, coursework, entrance exam scores, internships, and technical skills, through the centralized online job portal (Section 3.1.1). While early stages like application screening and written aptitude tests rely primarily on objective metrics (e.g., GPA cutoffs, written test scores), group debates and personal interviews introduce direct *social interactions*.

Group debates involve structured team discussions, allowing employers to assess communication, teamwork, and other socio-emotional skills. These interactions typically remain standardized—recruiters provide uniform prompts and ensure roughly equal speaking time per candidate—limiting opportunities to learn about individual backgrounds. In contrast, personal interviews often allow deeper, individualized social interactions related to candidates' interests, extracurricular activities, background, and interpersonal rapport (Deshpande, 2011; Rivera, 2012, 2015). These social interactions significantly enhance employers' insights into candidates' cultural fit and organizational alignment, making subtle indicators of social status more salient (Rivera, 2011; Rivera and Tilcsik, 2016; Laurison and Friedman, 2019). Although we lack data on interviewers' own backgrounds, ethnographic evidence shows that evaluative norms are quickly socialized within elite firms; thus managers with diverse origins likely converge on similar notions of "fit," reinforcing status-based biases (e.g., Turco, 2010; Bisin and Verdier, 2011; Castilla, 2015; Castilla and Ranganathan, 2020).

Section 4.6 below presents quantitative evidence on how interviews heighten the salience of (and indeed help reveal) social status, with complementary qualitative evidence in Online Appendix Section F.

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

When does caste become visible to employers? This subsection demonstrates quantitatively that personal interviews are where caste is revealed and penalized.

4.6.1 Quantitative Evidence

Although many surnames are caste-indicative, most do not reliably signal caste in urban settings such as elite colleges, where students come from diverse regions with varied naming conventions

(e.g., Jodhka, 2015; Parmar, 2020).[9] Consequently, only about 15% of students in my sample have clearly identifiable caste surnames. Figure 2 shows that, for this subgroup, caste earnings penalties appear earlier—at the application reading round. However, for the remaining 85% with non-identifiable surnames, the caste penalty emerges predominantly during personal interviews. The different timing of the caste penalty in the two subsamples strongly suggests that the caste penalty follows caste revelation. Thus, personal interviews serve as the key mechanism for revealing and penalizing social status.[10] Corresponding regression results, including robustness specifications, are shown in Online Appendix Tables OA.7 and OA.8.

Similar patterns in success rates. The same story demonstrated by Figure 2 holds when we look at success rates across screening rounds. Students whose surnames clearly reveal caste suffer large penalties immediately at the application reading round across almost every job (Online Appendix Figure OA.2). By contrast, students whose surnames are *not* caste-identifiable track the pattern shown for the overall sample in Table 6: their success rates are essentially equal to that of advantaged peers until the personal interview, where the penalty first appears and drives the final gap.

How are student names associated with caste? Surnames are coded as caste-identifiable when they appear in (i) the state-specific "schedules" that legally enumerate Scheduled Caste and Scheduled Tribe communities (hence the term *Scheduled*; see Puri, 2025), (ii) the *People of India* volumes compiled by the Anthropological Survey of India, which list common advantaged-caste surnames, or (iii) state-government gazette notifications that list disadvantaged communities (see Section 2.2) together with their commonly used surnames or titles.[11]

4.6.2 Qualitative Evidence

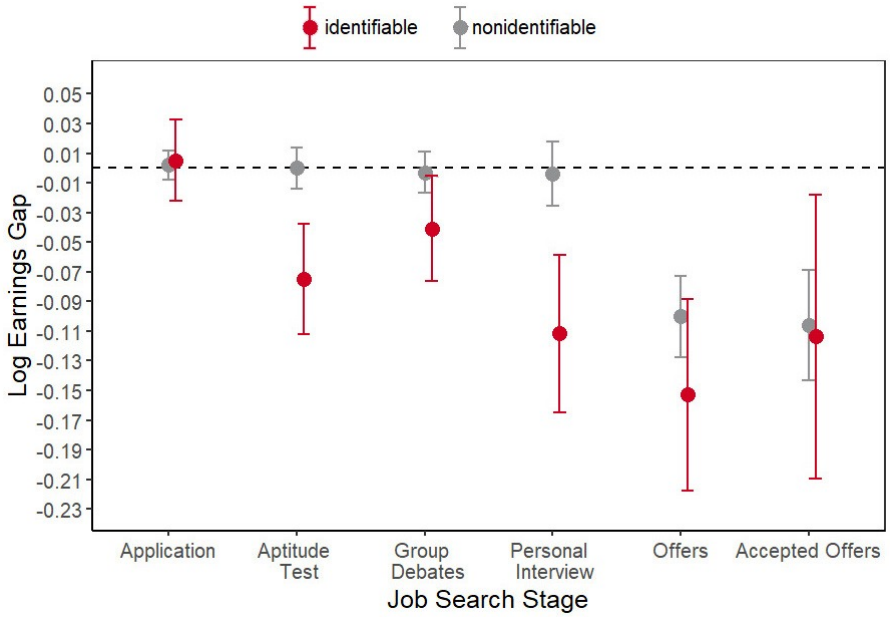
Qualitative evidence, consistent with the quantitative patterns above, indicates that: (i) visual and auditory cues like skin tone or regional accent do not reliably signal caste in urban, educated settings; and (ii) social interactions such as personal interviews often reveal social status in urban-educated contexts where other signals (e.g., names) are often ambiguous. See Online Appendix Section F for more details.

[9] See also Bhupatiraju et al. (2024)—a study of more than one million High Court cases from the urban centre of Patna—which shows that roughly half of petitioners, respondents, and judges use *caste-neutral* surnames (49% of petitioners, 56% of respondents, and 50% of judges), underscoring how common name-based ambiguity is in urban India.

[10] For example, a purely income-based explanation would predict an early gap in both groups, contrary to what we observe.

[11] The Election Commission of India publishes constituency-level voter rolls that flag seats reserved for Scheduled Castes or Scheduled Tribes; these rolls list the surnames or titles most commonly associated with the eligible disadvantaged groups in each area, providing an additional source for caste-linked surnames.

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



Notes: Figure 2 compares the log earnings gap across castes at each job search stage between students with identifiable (red) and non-identifiable (gray) caste surnames. Both gray and red dots show percentage differences in average salary between disadvantaged and advantaged castes, with 95% confidence intervals. Since salaries are identical and non-negotiable for identical jobs, gaps reflect differences in job access. All regressions include controls.

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

We examine *job-level heterogeneity* in the caste penalty to show that disadvantaged caste applicants experience about a 30 percentage point lower interview success rate across the *entire distribution* of jobs. Moreover, those who are hired exhibit roughly 20 percent higher promotion rates than their advantaged peers.

In the absence of screening round performance scores (Section 3.9), these headline findings raise two critical questions. First, might caste penalties appear earlier in *highly selective* jobs, which recruit primarily from the right tail of the ability distribution? Second, could personal interviews reveal genuine unobserved ability differences—rather than discrimination—that justify hiring disparities across castes? Below, I present two empirical facts followed by a model that exploits variation in job selectivity to distinguish ability differences from discrimination.

Fact 1 (Conditional Success Rates Reveal No Early Penalty, but a Large Final Interview Gap). Figure 5 plots, for each job, the conditional success rate gap (disadvantaged minus advantaged) against the cumulative cut rate at each screening stage—defined as the total fraction of applicants a job eliminates *by* the end of that stage. The cumulative cut rate proxies for *job selectivity*. The success rate gap for job $j \in \mathcal{J}$ is the coefficient β_j from the regression:

$$\Pr(\text{Success}_{ijt}^{\text{Job Search Stage}}) = \alpha_j + \beta_j \times \text{Disadv. Caste}_i + \text{Controls}_i + \gamma_t + \epsilon_{ijt}, \tag{2}$$

where $i \in \mathcal{I}$ indexes students and γ_t denotes cohort-year fixed effects. Controls include GPA, entrance exam scores, prior labor market experience, major, degree, and other characteristics.

Disadvantaged castes are about 30 percentage points less likely to advance than advantaged castes after personal interviews across all jobs. Before personal interviews, the caste gaps in advancement remain negligible, *even among highly selective jobs*. A distinct caste penalty emerges only after personal interviews. *This pattern holds even in uncontrolled regressions* (see Online Appendix Figure OA.3).

Fact 2 (Disadvantaged Hires Earn Comparably Higher Promotion Rates, Even in Highly Selective Jobs). I examine post-hire promotion rates using alumni records covering 98% of placed graduates (Section 3.8). I estimate caste differences in promotion rates using the following regression across job selectivity tiers:

$$\text{Outcome}_{ijt} = \alpha + \beta \times \text{Disadv. Caste}_i + \text{Controls}_i + k_j + \gamma_t + \epsilon_{ijt}, \tag{3}$$

where $i \in \mathcal{I}$ indexes students, $j \in \mathcal{J}$ indexes jobs, γ_t denotes year fixed effects, and k_j represents job fixed effects. Controls include GPA, entrance exam scores, prior labor market experience, major, degree, and other characteristics.

Table 7 presents these results. Pooling all jobs (columns 1-2), disadvantaged hires have approximately 20% higher promotion rates than advantaged peers. Crucially, when jobs are grouped by selectivity tiers based on cumulative cut rates (columns 3-9), this promotion advantage remains stable at 20-25% across the selectivity distribution. Even in the most selective jobs (top 10%), the point estimate remains comparable though less precise. The pattern holds for both client-facing and non-client-facing roles (columns 10-11). The high baseline promotion rates (around 0.70) create boundary bias that inherently limits the scope for unobservables to alter estimates.[12] Indeed, Oster (2019) bounds confirm this, showing minimal coefficient changes (see Online Appendix Section E).

Promotions as a performance metric. In entry-level jobs at multinationals, pay is often globally fixed at company headquarters, making promotions the main channel for rewarding productivity and thus a reliable performance proxy (e.g., Bloom et al., 2012; Minni, 2024a).

[12]The baseline promotion rates increase modestly with firm selectivity. This pattern aligns with more selective firms typically featuring stronger internal labor markets and promotion-based incentive structures (Doeringer and Piore, 1971; Baker et al., 1994), tournament-style advancement systems (Lazear and Rosen, 1981), and up-or-out cultures particularly in elite professional services (Kahn and Huberman, 1988). The higher baseline rates at selective firms—which invest heavily in employee development having already screened for quality—make the persistent 20-25% caste promotion advantage even more striking.

Promotion rate estimates are conservative. Online Appendix Table OA.9 shows that job leaving entails little GPA-based selection within a given job. Before controlling for job fixed effects, the leaver–stayer GPA penalty is larger for advantaged castes than for disadvantaged castes. Two additional pieces of context suggest promotion estimates are likely conservative. First, promotion decisions inevitably involve greater managerial discretion, and such subjective assessments typically favor higher-status candidates (e.g., Turco, 2010; Rivera, 2015; Castilla, 2015; Benson et al., 2019, 2023). Second, caste or even broader class-based initiatives are virtually absent at top firms—fewer than 10% of major U.S. firms reference social class in DEI statements (Stansbury and Rodriguez, 2024b), and even Indian multinationals seldom mention caste (Bapuji et al., 2023).

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

The interview penalty (Fact 1) followed by the subsequent promotion advantage (Fact 2) could simply reflect more right-tail performers among the hired disadvantaged (inframarginality) instead of stricter hiring thresholds. However, that the promotion advantage is *positive and stable* across all job selectivity tiers suggests discrimination rather than unobserved ability differences. I formalize this intuition through a model demonstrating how (i) the absence of early-stage gaps for disadvantaged castes implies no pre-interview ability deficit, and (ii) variation in screening selectivity across jobs helps distinguish discrimination from unobserved ability differences revealed at interviews.

Consider jobs $j \in \mathcal{J}$ with four sequential screening rounds $s \in \{1, 2, 3, 4\}$, where $s = 1, 2, 3$ are pre-interview stages (application, written tests, group debates) and $s = 4$ is the personal interview. Two caste groups $g \in \{0, 1\}$, where $g = 1$ denotes disadvantaged castes and $g = 0$ denotes advantaged castes, draw screening scores $M_{s,i,g}$ from distributions $F_{s,g}$, common across jobs within each stage. Using residualized score distributions leaves all results unchanged.

Jobs set thresholds $\tau_{s,j}$ for each round, but crucially, these may differ by caste only at the interview stage, $\tau_{4,j}^{(g)}$, reflecting that caste becomes salient to the employer at this stage (Section 4.6). A candidate passes round s iff their score satisfies $M_{s,i,g} \geq \tau_{s,j}$ and they have passed all preceding rounds. A caste penalty at interviews is indicative of discrimination if $\tau_{4,j}^{(1)} > \tau_{4,j}^{(0)}$. I make five main assumptions:

Assumption 1 (**Perfect information**)**.** Ability is perfectly proxied by the screening scores, and employers observe those scores without error, at each stage. This simplifies exposition but is not essential, as our main results hold under imperfect information (Online Appendix Section G).

Assumption 2 (**Regularity Conditions**)**.** The performance distributions $F_{s,g}$ are continuous with smooth densities and finite moments.

Assumption 3 (**Monotonicity of thresholds**)**.** We assume that thresholds are monotonically increasing across screening rounds:

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

Assumption 4 (**Common Talent Pool Across Jobs**)**.** All jobs recruit from the same underlying talent pool: $F_{s,g}$ does not vary across jobs j. This assumption is reasonable because the jobs students apply to are similar across castes (Figure 1) and there are negligible caste gaps in advancement before interviews across job selectivity tiers (Fact 1).

Assumption 5 (**Light Tails**)**.** For each caste g and stage s, let the *hazard rate* be $h_{s,g}(m) \equiv f_{s,g}(m)/\big[1 - F_{s,g}(m)\big]$. There exist constants $C_g > 0$ and $\beta > 0$ such that

$$\Pr\big(M_{s,i,g} \geq m\big) \;\leq\; C_g\, e^{-\beta m}, \qquad h_{s,g}(m) \xrightarrow[m\to\infty]{} \infty.$$

Since screening round scores are bounded in practice (e.g., test scores between 0-100), the distributions have finite support, trivially satisfying this light-tail condition.

Proposition 1 (**No Early Gap Implies No Right-Tail Ability Deficit Before Interviews**)**.** Consider a job j with four rounds of screening, with $s = 1, 2, 3$ the pre-interview stages and $s = 4$ the personal interview. Under Assumptions 1–5, if disadvantaged castes ($g = 1$) were genuinely weaker in the right tail of the ability distribution, then for a sufficiently high threshold, $\tau_{s,j}$ at any pre-interview stage $s < 4$, the within-job pass-rate gap

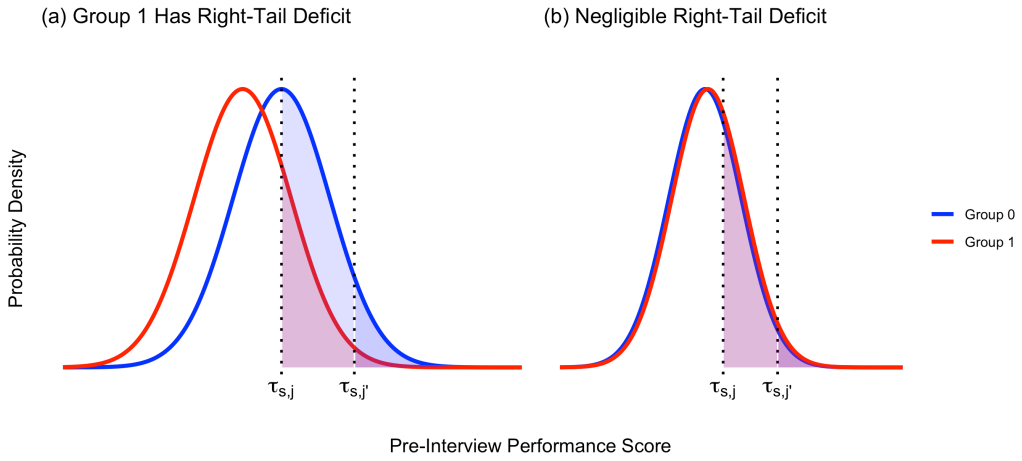
$$\Delta_{s,j} \equiv P(\text{pass}_{s,j} \mid g = 1) - P(\text{pass}_{s,j} \mid g = 0)$$

would become negative. Conversely, observing $\Delta_{s,j} \geq 0$ for jobs, including the most selective, rules out any meaningful right-tail ability deficit for disadvantaged castes before interviews.

Sketch of Proof. Figure 3 shows why variation in job selectivity reveals ability differences before interviews. Under Assumption 4, cross-job variation in the pre-interview cutoffs $\tau_{s,j}$ provides *exogenous* shifts that probe progressively deeper into both groups' *common* right tail.

- **Panel (a): Right-tail deficit.** If disadvantaged castes had weaker right-tail performance before interviews, then by Assumption 5, their pass rates would decline *faster* than advantaged castes' as $\tau_{s,j}$ increases. Moreover, this divergence would become increasingly apparent at higher cutoffs, eventually yielding $\Delta_{s,j} < 0$ for sufficiently selective jobs.
- **Panel (b): No right-tail deficit.** Fact 1 shows $\Delta_{s,j} \approx 0$ across all jobs, even the most selective, ruling out meaningful right-tail performance deficits for disadvantaged castes before interviews.

Figure 3: **Testing for Pre-Interview Performance Differences**



Notes: Group 0 denotes advantaged castes; Group 1 denotes disadvantaged castes. $\tau_{s,j}$ denotes screening threshold at stage $s < 4$ for job j (j' represents a more selective job). Shaded regions show candidates who pass the thresholds (blue for Group 0, red for Group 1, purple where they overlap). Panel (a): A hypothetical right-tail deficit for Group 1 would become more apparent at higher cutoffs. Panel (b): Ability distributions implied by Fact 1 show no right-tail deficit before interviews.

□

Do interviews reveal unobserved ability differences that justify disparities? Given Proposition 1 and the differential timing of the caste penalty by subsample (Figure 2), any observed late-stage gap at the interview stage must stem either from newly revealed ability signals during interviews or from discrimination based on social status. To distinguish between these explanations, I examine how post-hire promotion patterns vary with job selectivity.

Let job performance be $\bar{M}_{i,g}$, the average of a candidate's four screening scores. Candidate i is promoted iff $\bar{M}_{i,g} \geq \kappa_j$, where κ_j is a job-specific promotion cutoff. For jobs in selectivity tier q, let $\tau^q_{4,j}$ denote the common interview cutoff that would apply to both castes in the top $(1-q)$ fraction of jobs by selectivity (e.g., $q = 0.8$ means the top 20% most selective jobs). Let $\tau^{(g)}_{4,j}$ denote the actual group-specific cutoff that caste g faces in job j. Under no discrimination, $\tau^{(1)}_{4,j} = \tau^{(0)}_{4,j} = \tau^q_{4,j}$; under discrimination, $\tau^{(1)}_{4,j} > \tau^{(0)}_{4,j}$. Define the average screening score of job j's hires from caste g as:

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

and the promotion gap in selectivity tier q as $\Delta^{\text{Prom}}(q) = \bar{Y}^{(1)}_j - \bar{Y}^{(0)}_j$.

Proposition 2 (Stable Promotion Gaps Across Selectivity Tiers Indicate Discrimination). Fix a selectivity tier q. Let $\tau^q_{4,j}$, $\tau^{(g)}_{4,j}$, $\bar{Y}^{(g)}_j$, and $\Delta^{\text{Prom}}(q)$ be defined as above. Under Assumptions 1–5:

1. *No discrimination.* If $\tau^{(1)}_{4,j} = \tau^{(0)}_{4,j} = \tau^q_{4,j}$ for all jobs in tier q, then as selectivity increases:

$$\Delta^{\text{Prom}}(q) \to 0 \quad \text{as } q \to 1$$

2. *Discrimination.* If $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} + \delta$ with $\delta > 0$, then:

$$\Delta^{\text{Prom}}(q) \to \delta > 0 \quad \text{as } q \to 1$$

Thus, a positive and stable promotion gap as job selectivity increases indicates discrimination.

Sketch of Proof. Under Assumption 4, which makes cross-job variation in $\tau_{4,j}$ plausibly exogenous, the interview cutoff acts as an *instrument*: it varies across jobs and affects who gets hired (relevance) while affecting promotions only through selection (exclusion). The key insight is that as interview cutoffs $\tau_{4,j}$ rise across increasingly selective jobs, *the average ability of hires converges to that of the marginal hire.*

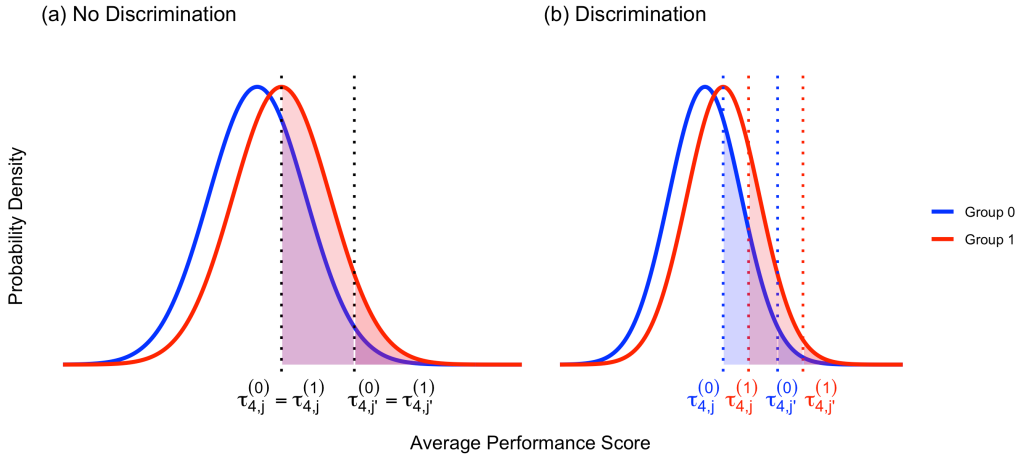
- **Panel (a): No discrimination.** If both castes face $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} = \tau$, then with light tails (Assumption 5):

$$\bar{Y}_j^{(g)} = \tau + \frac{1}{h(\tau)}, \quad \text{where } h(\tau) = \frac{f(\tau)}{1 - F(\tau)} \xrightarrow[\tau \to \infty]{} \infty$$

  Hence $\bar{Y}_j^{(g)} - \tau \to 0$ for both groups, so $\Delta^{\text{Prom}}(q) \to 0$ as selectivity increases (complete derivation in Online Appendix Section H).

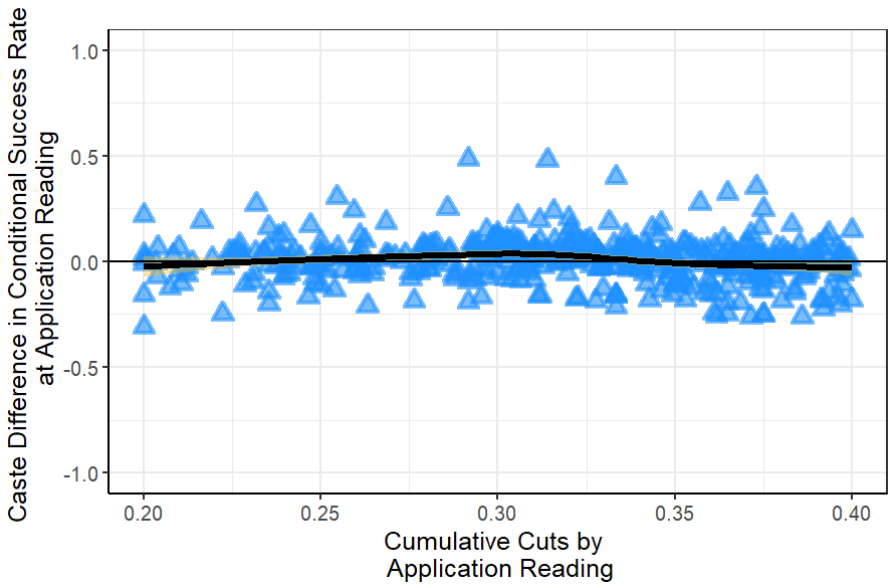
- **Panel (b): Discrimination.** If disadvantaged castes face $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} + \delta$ with $\delta > 0$, then $\bar{Y}_j^{(1)} - \bar{Y}_j^{(0)} \to \delta > 0$ regardless of how high cutoffs rise. Thus, a positive and stable promotion advantage persists even in the most selective jobs, consistent with $\delta > 0$.

Figure 4: **Testing for Discrimination**

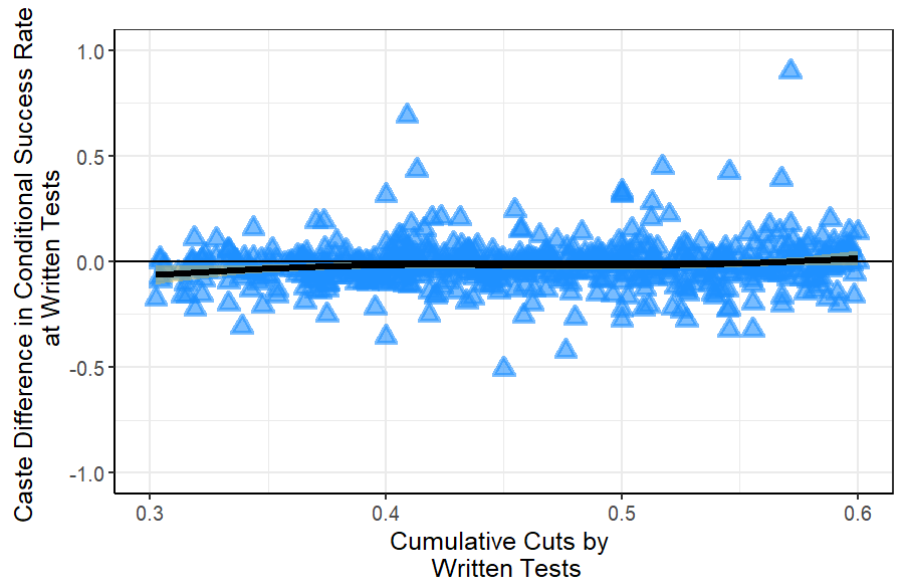


Notes: Group 0 denotes advantaged castes; Group 1 denotes disadvantaged castes. $\tau_{4,j}^{(g)}$ denotes group g's interview threshold; j' represents a more selective job than j. Shaded regions show candidates who pass the thresholds (blue for Group 0, red for Group 1, purple where they overlap). Panel (a): Equal cutoffs lead to vanishing performance gaps as selectivity increases. Panel (b): A constant cutoff gap $\delta > 0$ maintains a stable performance advantage for Group 1. □

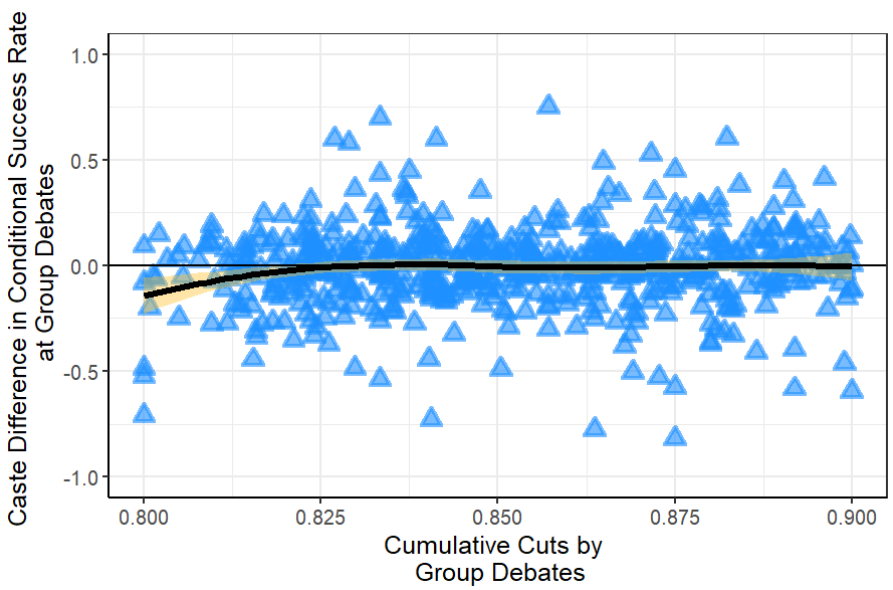
Figure 5: **Caste Penalty in the Conditional Success Rate At Each Screening Round**



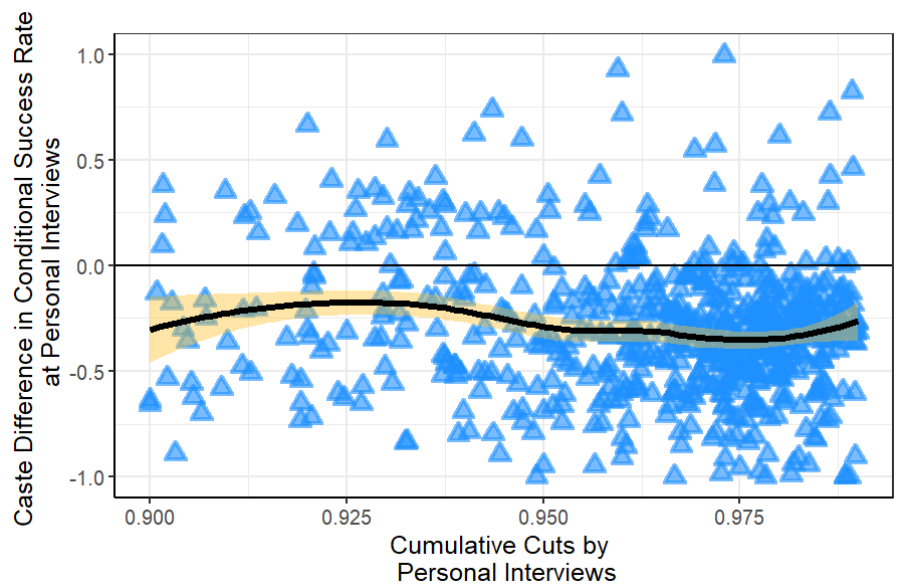
(a) Caste Penalty in Conditional Success Rates versus Cumulative Cuts by Application Reading



(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 plots caste differences in conditional success rates against cumulative cuts at each screening round. Negative values indicate lower advancement rates for disadvantaged castes relative to advantaged castes. Blue triangles represent individual jobs, with smoothed conditional means (black lines) and 95% confidence intervals (yellow). The x-axis shows the cumulative proportion of students eliminated by each round; the y-axis shows the caste coefficient from job-specific regressions with controls (degree, major, GPA, entrance exam scores) and cohort-year fixed effects. The horizontal line at zero indicates no caste difference.

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 Top 10%	Promoted at Client Facing	Promoted at Non-Client Facing
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Disadv. Caste (OLS)	0.137***	0.096***	0.104***	0.112***	0.136***	0.147***	0.190**	0.204**	0.084	0.107**	0.082**
	(0.025)	(0.026)	(0.029)	(0.038)	(0.043)	(0.054)	(0.073)	(0.084)	(0.109)	(0.044)	(0.032)
Tenure Duration		0.084***	0.082***	0.076***	0.074***	0.068***	0.054*	0.079**	0.115**	0.065***	0.097***
		(0.012)	(0.014)	(0.019)	(0.021)	(0.025)	(0.030)	(0.037)	(0.049)	(0.021)	(0.015)
Mean Y	0.684	0.684	0.667	0.668	0.676	0.711	0.711	0.706	0.730	0.687	0.682
Observations	2,489	2,489	2,012	1,246	997	748	502	374	252	835	1,654
Adjusted R^2	0.183	0.205	0.203	0.216	0.211	0.182	0.192	0.189	0.127	0.194	0.214
Job FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Other Controls		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Notes: Table 7 reports caste differences in promotion rates for graduates who obtained jobs through the main placement process. Column (1) includes only job and year fixed effects. Column (2) adds controls for GPA, major, degree, and job tenure. Columns (3)-(9) show results by job selectivity tiers, measured by cumulative applicant elimination rates. Columns (10)-(11) compare client-facing versus non-client-facing roles. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

4.8 Classification and Preferred Interpretation of the Caste Penalty: Distinguishing from Statistical or Customer Discrimination

The most plausible interpretation for caste hiring disparities is *disparate treatment* (explicit caste-based discrimination) driven by *bias*. Four facts jointly support this view. First, the caste penalty closely follows caste revelation (Section 4.6). Second, disadvantaged castes face not only lower hiring rates but penalties that are remarkably consistent across the *entire* job distribution (Fact 1). Third, disadvantaged hires have persistently higher and stable promotion rates across all job selectivity tiers (Fact 2). Fourth, disadvantaged caste hires are more likely to be promoted in both client-facing and non-client-facing jobs (Table 7). Taken together, these findings are most consistent with bias rather than accurate statistical inference about productivity[13] or customer-driven discrimination. Although subjective criteria for assessing "fit" might inadvertently create a disparate impact (e.g., by favoring candidates with particular extracurricular backgrounds), the cumulative evidence presented in Sections 4.2–4.7 points most strongly to bias-driven disparate treatment.

4.9 What Happens During Personal Interviews?

We do not directly observe interview transcripts, so we cannot parse the exact conversations. However, prior ethnographic research on elite hiring demonstrates that personal interviews place substantial weight on subjective markers of "fit" or "cultural match" that are closely tied to candidates' socioeconomic backgrounds rather than strictly job-relevant skills (e.g., Rivera, 2012, 2015; Lamont, 2009). In Rivera's comprehensive study of U.S. professional service firms, more than half of all evaluators ranked "cultural fit" as the single most important hiring criterion—ahead of analytical thinking and communication skills. This emphasis varies by industry: approximately 70% of evaluators in law firms, 55% in investment banking, and 40% in consulting cited "fit" as their top priority. Strikingly, roughly four-fifths of interviewers employed some variant of the so-called "airport test"—asking themselves whether they would want to be stuck in an airport with the candidate for several hours. As one investment banking director explained:

> "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?" Rivera (2012, p. 1010).

A richer qualitative account of how interviewers lean on cultural cues during such social interactions, complete with illustrative quotations, is provided in Online Appendix Section I.

[13] An accurate statistical discrimination model typically assumes evaluators hold correct beliefs about average group-based productivity and that the underlying productivity distributions are relatively well-behaved (Aigner and Cain, 1977).

5 Alternative Explanations

Below, I briefly consider and rule out alternative explanations for the observed caste disparities emerging specifically during personal interviews.

1. **Ability gaps.** If disadvantaged candidates truly lagged in ability—either on cumulative performance across pre-interview rounds or on specific technical and socio-emotional skills needed for these roles—such deficits would emerge earlier or most sharply in the most selective jobs. Yet Fact 1 shows no statistically significant caste penalty in pre-interview rounds across all selectivity tiers. Moreover, Fact 2 shows hired disadvantaged candidates subsequently outperform advantaged peers across all jobs.

2. **Delayed discrimination (strategic optics).** Employers have little incentive to delay caste discrimination: indeed, when caste is identifiable through surnames, penalties emerge immediately at the application reading round (Figure 2; Online Appendix Figure OA.2).

   Relatedly, there is weak institutional pressure on private-sector employers to promote caste diversity. Unlike the U.S., India lacks a federal ombudsman such as the Equal Employment Opportunity Commission to regulate private-sector hiring. Furthermore, caste, class, or other forms of socioeconomic status indicators are *not* diversity metrics that multinational firms typically monitor, with race and gender being the main focus areas for diversity in such firms (Stansbury and Rodriguez, 2024a,b). Additionally, only 3% of India's top 100 listed firms collected caste data for HR records in 2018 (BusinessLine, 2018), with many Indian multinational companies explicitly viewing caste-targeted hiring as "anti-merit" (Jodhka and Naudet, 2019). Even as recently as 2023, most firms seldom address caste discrimination in their diversity initiatives (Bhupatiraju et al., 2024). Hence, even absent formal government pressure, internal institutional pressure remains weak.

3. **Attrition within standardized job ladders.** Our promotion regressions compare workers who start in the *same* role at the *same* firm, and promotion schedules for those roles are highly centralized across multinational subsidiaries (e.g., Bloom et al., 2012). Additionally, as mentioned in Section 4.7, promotion rate estimates are likely conservative.

4. **Preferential promotion for disadvantaged hires.** Given the rarity of caste-focused corporate initiatives mentioned above, and extensive survey evidence consistent with widespread bias *against* disadvantaged caste employees rather than in their favor (e.g., California v. Cisco Systems, Inc., 2022), it is unlikely that the promotion results are confounded by policies explicitly favoring disadvantaged castes at multinationals.

5. **Missing data or skipped rounds.** For students participating in the placement process, the placement office enforces mandatory attendance at every stage and strictly prohibits employers from skipping steps or altering recruitment procedures mid-process (Section 3.1.2). Comprehensive exit surveys conducted by the placement office independently corroborate the completeness of official records (Section 4.1).

6. **Sample attrition and external search.** Once students register and submit applications through the centralized portal, virtually no one exits the placement process for non-elimination reasons. Exits occur only through formal elimination decisions at each of the four strictly enforced screening rounds, with the placement office meticulously recording every advance-or-exit decision (Sections 3.1.1 and 3.1.2). Return offers from internships are rare (Section 2.3), and the placement office prohibits registered candidates from conducting parallel off-campus searches[14], leaving little scope for differential outside options to drive the observed gap.

7. **Negotiation and salary differences.** Salaries and benefits are ex-ante standardized, strictly non-negotiable during the course of the job fair (Section 2.3; Section 3.4), and verified by the placement office through comparing offer letters with the compensation schedules advertised by firms (Sections 2.3 and 4.1).

8. **Public-sector competition.** Government jobs have caste-specific quotas, but they account for only 2% of all campus job listings (Table 3) and are dropped from the main analysis (Section 3.4). Additionally, all job listings are visible to both candidates and recruiters through a centralized placement portal (Section 3.1.1), ensuring private-sector employers are fully aware of the limited public-sector competition.

9. **Preferences over jobs.** Near-zero application costs (Section 3.6) lead students to apply broadly to the same set of jobs via the centralized portal (Section 3.1.1). In addition, caste differences in preferences over amenities at best explain only a negligible share of the initial caste earnings penalty (Figure 1).

6 External Validity

How generalizable are findings from one elite Indian institution? While caste discrimination might initially appear uniquely Indian, five considerations underscore why these results hold broader implications, not only across India but globally in elite hiring and other high-stakes selection contexts.

[14] For common placement office rules at elite Indian colleges, see e.g., Central Placement Cell, Delhi University.

1. Mirroring global elite hiring pipelines. The hiring process we analyze—résumé reading, preliminary screens (e.g., written tests), and discretionary final interviews—is not only replicated across the job fairs of elite Indian colleges (Section 2.4) but is also standard screening procedure at elite jobs worldwide (Rivera, 2015, 2012). Indeed, almost all high-paying jobs use a combination of objective and subjective assessments in their screening (e.g., Highhouse, 2008; Laurison and Friedman, 2019). Moreover, multinational corporations often use standardized evaluation protocols across their global offices (e.g., Bloom et al., 2012; Minni, 2024a). By pinpointing class bias at the final interview stage, our evidence speaks directly to mainstream hiring at global consultancies, tech firms, and banks (Rivera, 2011).

2. Subjective "fit" screens reproduce class disparities worldwide. Class disparities that arise at subjective gatekeeper stages are a structural feature of elite selection globally. Across settings, discretionary assessments tend to reward higher-status candidates' cultural capital (Bourdieu, 1986; Rivera, 2015). As explicit status markers fade in diverse urban labor markets (Clark, 2014; Jodhka, 2015; Rose, 2023), these subjective evaluations become critical gatekeepers. Evidence from U.S. and U.K. professional hiring shows that brief, rapport-heavy interviews engender class gaps by privileging mannerisms, conversational style, and extracurriculars that signal organizational fit (Rivera, 2012; Laurison and Friedman, 2019). Similarly, in U.S. academia, class influences who receives NSF grants and secures tenure at top research institutions (Stansbury and Rodriguez, 2024a). U.S. college admissions reveal parallel dynamics, as "holistic" file reads involving personal essays, recommendation letters, extracurricular narratives, and alumni interviews systematically disadvantage applicants from lower socioeconomic backgrounds (Borneman et al., 2007; Kuncel et al., 2013; Alvero et al., 2021; Sacerdote et al., 2025).

India's elite civil service provides striking cross-national validation of these patterns: disadvantaged caste officers who scored high enough in their combined written and subjective assessments to qualify without quotas, yet received markedly lower ratings in their "personality" interviews subsequently outperform advantaged peers by 18–20% (Bhavnani and Lee, 2021).

3. Downstream erosion of upward-mobility policies. Policies widening entry at the *front end* of elite pipelines often lose traction *downstream*. In India, caste-based admissions quotas boost elite college representation, yet discretionary hiring screens at multinationals erode these gains, as this study demonstrates. Similar patterns emerge globally: in the U.S., need-blind admissions and affirmative action expand educational access, yet high-income graduates are far more likely to convert prestigious college degrees into elite careers through exclusive social networks (Michelman et al., 2022). In Chile, students from lower socioeconomic backgrounds

face persistent gaps in reaching leadership roles despite attending top colleges (Zimmerman, 2019). These patterns reveal that admissions reforms alone are insufficient in advancing upward mobility for disadvantaged groups.

4. A globally competitive talent pool. The students in our sample match their Ivy-Plus counterparts in human-capital quality. Using harmonized alumni earnings metrics for nearly 2,800 institutions around the world, Martellini et al. (2024) place several elite Indian universities alongside leading global universities, including many Ivy League institutions. Hence, the labor-market segment we study is on par with the world's most talented.

5. Technically-oriented roles as a conservative test of class discrimination. Our focus on technical positions—software engineering, data analytics, quantitative consulting—likely represents a *best-case* scenario for meritocratic hiring. These roles prioritize demonstrable technical skills, use standardized written or coding tests, and rely on objective performance metrics (Graham, 2004), theoretically leaving minimal room for social status considerations (Bourdieu, 1986) and thus likely providing a conservative test of class-based hiring discrimination. Technically-oriented cultures often pride themselves on valuing competence over status, with industry lore celebrating engineering prowess regardless of background (Turkle, 1984; Ensmenger, 2010). That we find substantial class penalties even in this supposedly meritocratic domain is particularly striking, especially given recent evidence that subjective criteria like "potential" poorly predict performance *even when* firms could plausibly claim they are job-relevant (Benson et al., 2023).

Taken together, these five considerations suggest that the interview-stage caste penalty we document is not an idiosyncrasy of one Indian campus but a vivid instance of a broader pattern: whenever elite selection relies on discretionary "fit" assessments, class disparities often resurface, eroding the promise of meritocracy across contexts and countries.

7 Conclusion

Personal interviews for elite jobs afford recruiters the opportunity to discriminate on social status through subjective "fit" assessments. By analyzing how leading multinationals recruit in India, I provide the first large-scale evidence on where and how class-based discrimination emerges in the elite hiring pipeline as well as its mechanisms. Three findings stand out: First, nearly 90% of caste-based disparities arise during subjective personal interviews, not during relatively more standardized assessments. Second, hired disadvantaged castes have about 20%

higher promotion rates than advantaged peers across all roles. Third, the hiring penalty is most consistent with bias.

The Indian setting provides unusual clarity for investigating social status discrimination. Although caste is uniquely Indian, the underlying mechanism—subjective evaluations of who "belongs"—operates globally. Class discrimination often manifests through social interactions revealing cues like conversational style, cultural pursuits, and shared experiences (Rivera, 2015; Laurison and Friedman, 2019). Dual-process cognition theory provides psychological underpinnings for this pattern: when facing ambiguous hiring decisions, evaluators typically default to rapid, intuitive judgments about cultural similarity rather than slow, deliberative assessments of competence (Kahneman, 2011). These snap impressions, shaped by superficial rapport (Todorov, 2017), implicit stereotypes (Greenwald and Banaji, 1995), and immediate affective reactions (Pinker, 2021), typically predict performance no better than chance, yet are systematically biased against lower social status candidates (Highhouse, 2008; Dana et al., 2013; Rivera, 2015).

The problem extends beyond individual bias to structural organizational incentives. Hiring managers at multinationals rarely supervise the candidates they select (e.g., Cappelli, 2019), creating an agency problem where decision-makers bear little consequence for poor choices (Hoffman et al., 2018). At elite firms, where reaching final interviews signals strong qualifications, subjective fit assessments disproportionately exclude candidates from "unfamiliar" backgrounds (Rivera, 2012; Laurison and Friedman, 2019), reflecting misaligned incentives that reward cultural homogeneity over performance potential (Hoffman and Stanton, 2024; Benson et al., 2023).

These findings expose a critical flaw in diversity efforts: gains from affirmative action in education can erode when graduates face subjective hiring screens that provide recruiters leeway to penalize social status. Even as elite universities diversify, employers could restore homogeneity by selecting for cultural fit (Rivera, 2012, 2015). This pattern reflects deeper dynamics of cultural transmission, where elite practices co-evolve to maintain class boundaries (Bisin and Verdier, 1998, 2024, 2025). When hiring managers select for familiar cultural signals, they perpetuate class advantages that prove remarkably persistent across generations, even in the face of policy interventions designed to promote mobility (Barrios-Fernández et al., 2024).

Our results do *not* imply that all interviews are doomed to bias. Well-structured interviews—with job-relevant questions asked in fixed order and scored against explicit rubrics—can often match the predictive validity of other standardized assessments while reducing demographic disparities (e.g., Campion et al., 1997; Levashina et al., 2014). Context matters: small teams hiring for roles they directly supervise tend to weigh task competence over social similarity (Hinds et al., 2000), whereas HR gatekeepers who rarely work with new hires have greater latitude to default to fit-focused heuristics (Rivera and Tilcsik, 2016). While some cultural

similarity can facilitate coordination (Reagans and Zuckerman, 2001), it becomes costly when such signals crowd out more relevant indicators of capability (Gorman, 2005; Castilla, 2015). Even if screening on "fit" yields efficiency gains for individual firms, concentrating elite opportunities among the already privileged forgoes broader gains from diversifying leadership (Page, 2007; Chetty et al., 2018).

This study captures only entry-level hiring and initial career outcomes. Future research should track whether performance advantages persist long-term or whether early discrimination compounds through network effects, potentially reversing these gains through cumulative disadvantages in social capital, sponsorship, and access to high-visibility opportunities. Moreover, while we observe hiring outcomes, we cannot see interview content itself—the specific psychological triggers that activate bias. Analyzing actual interview interactions would illuminate the psychological mechanisms through which class markers influence evaluator judgments (Kraus et al., 2017). Given the role of quick judgments in facilitating bias (Kahneman, 2011), experimental work could test interventions that disrupt automatic evaluations—from structured interview formats to forced deliberation pauses to third-party review mechanisms that prompt analytical thinking (e.g., Stanovich and West, 2000; Tetlock, 2005; Tetlock and Mellers, 2011).

What can organizations do? While eliminating interviews is impractical, structured behavioral questions can reduce bias while preserving valid assessment (Highhouse, 2008). Most critically, tracking interviewer ratings against subsequent performance creates accountability—when evaluators' "fit" assessments are audited against productivity-relevant outcomes, bias is often mitigated (Mero and Motowidlo, 1995; Castilla, 2015). Whether in courts (Albright, 2019), admissions (DeVaul et al., 1987), or employment (Rivera, 2012), whenever discretion enters ostensibly neutral processes, class advantages often reassert themselves. Breaking this cycle requires recognizing that our intuitions about who "belongs" often reflect inherited privilege rather than individual merit. Only by constraining these impulses—through structure, accountability, and continuous monitoring—can organizations tap overlooked talent and help societies achieve genuine mobility.

References

[1] Ægisdóttir, S., M. J. White, P. M. Spengler, A. S. Maugherman, L. A. Anderson, R. S. Cook, C. N. Nichols, G. K. Lampropoulos, B. S. Walker, G. Cohen, and J. D. Rush (2006). The meta-analysis of clinical judgment project: Fifty-six years of accumulated research on clinical versus statistical prediction. *The Counseling Psychologist 34*(3), 341–382.

[2] Aigner, D. J. and G. G. Cain (1977). Statistical theories of discrimination in labor markets. *Industrial and Labor Relations Review 30*(2), 175–187.

[3] Albright, A. (2019). If you give a judge a risk score: Evidence from Kentucky bail decisions. Working paper.

[4] 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.

[5] Alvero, A., S. Giebel, B. Gebre-Medhin, A. L. Antonio, M. L. Stevens, and B. W. Domingue (2021). Essay content and style are strongly related to household income and SAT scores: Evidence from 60,000 undergraduate applications. *Science Advances 7*(42), eabi9031.

[6] Baker, G., M. Gibbs, and B. Holmström (1994). The internal economics of the firm: Evidence from personnel data. *Quarterly Journal of Economics 109*(4), 881–919.

[7] Bapuji, H., K. Gupta, S. Chrispal, and T. Roulet (2023). What managers everywhere must know about caste. *MIT Sloan Management Review Winter*. Accessed: 2024-04-22.

[8] Barrios-Fernández, A., C. Neilson, and S. Zimmerman (2024). Elite universities and the intergenerational transmission of human and social capital. *American Economic Review*. Conditionally accepted.

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

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

[11] Benson, A., D. Li, and K. Shue (2023). "Potential" and the gender promotion gap. Conditionally accepted.

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

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

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

[15] 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.

[16] Bhupatiraju, S., D. L. Chen, S. Joshi, and P. Neis (2024). Caste aside? Names, networks and justice in the courts of Bihar, India. *European Journal of Empirical Legal Studies 1*(2), 151–178.

[17] Bisin, A., J. Rubin, A. Seror, and T. Verdier (2023). Culture, institutions & the long divergence. *Journal of Economic Growth 29*, 1–40.

[18] Bisin, A. and T. Verdier (1998). On the cultural transmission of preferences for social status. *Journal of Public Economics 70*(1), 75–97.

[19] Bisin, A. and T. Verdier (2011). The economics of cultural transmission and socialization. In *Handbook of Social Economics*, Volume 1, pp. 339–416. Elsevier.

[20] Bisin, A. and T. Verdier (2024). On the joint evolution of culture and political institutions: Elites and civil society. *Journal of Political Economy 132*(5).

[21] Bisin, A. and T. Verdier (2025). Economic models of cultural transmission. CEPR Discussion Paper DP20340, CEPR Press, Paris and London.

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

[23] Bol, T., J. Witschge, H. G. Van de Werfhorst, and J. Dronkers (2014). Curricular tracking and central examinations: Counterbalancing the impact of social background on student achievement in 36 countries. *Social Forces 92*(4), 1545–1572.

[24] Borneman, M. J., S. R. Cooper, D. M. Klieger, N. R. Kuncel, and N. Brenckman (2007). Admission interviews: Not so nice to meet you. In *Alternative Predictors of Academic Performance, Symposium conducted at the Annual Meeting of the Society for Industrial and Organizational Psychology, New York, NY.*

[25] 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.

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

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

[28] Campion, M. A., D. K. Palmer, and J. E. Campion (1997). A review of structure in the selection interview. *Personnel Psychology 50*(3), 655–702.

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

[30] Castilla, E. J. (2015). Accounting for the gap: A firm study manipulating organizational accountability and transparency in pay decisions. *Organization Science 26*(2), 311–333.

[31] Castilla, E. J. and S. Benard (2010). The paradox of meritocracy in organizations. *Administrative Science Quarterly 55*(4), 543–576.

[32] Castilla, E. J. and A. Ranganathan (2020). The production of merit: How managers understand and apply merit in the workplace. *Organization Science 31*(4), 797–1051.

[33] 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.

[34] Chetty, R., N. Hendren, M. R. Jones, and S. R. Porter (2018). Race and economic opportunity in the United States: An intergenerational perspective. *The Quarterly Journal of Economics 135*(2), 711–783.

[35] Chow, R. M. and E. D. Knowles (2016). Taking race off the table: Agenda setting and support for color-blind public policy. *Personality and Social Psychology Bulletin 42*(1), 25–39.

[36] Clark, G. (2014). *The Son Also Rises: Surnames and the History of Social Mobility*. The Princeton Economic History of the Western World. Princeton University Press.

[37] Comerford, D., J. V. Rodríguez Mora, and M. J. Watts (2022). Meritocracy and the inheritance of advantage. *Journal of Economic Growth 27*(2), 235–272.

[38] Dana, J., R. Dawes, and N. Peterson (2013). Belief in the unstructured interview: The persistence of an illusion. *Judgment and Decision Making 8*(5), 512–520.

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

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

[41] DeVaul, R., F. Jervey, J. Chappell, P. Caver, B. Short, and S. O'Keefe (1987). Medical school performance of initially rejected students. *Journal of the American Medical Association 257*, 47–51.

[42] Doeringer, P. B. and M. J. Piore (1971). *Internal Labor Markets and Manpower Analysis*. Lexington, MA: Heath Lexington Books.

[43] Durlauf, S. N. (1999, Fall). The case "against" social capital. *Focus 20*(3), 1–5.

[44] Durlauf, S. N. (2006). Groups, social influences, and inequality: A memberships theory perspective on poverty traps. In S. Bowles, S. N. Durlauf, and K. Hoff (Eds.), *Poverty Traps*. Princeton: Princeton University Press.

[45] Durlauf, S. N. (2008). Memberships and inequality. In D. Castiglione, J. van Deth, and G. Wolleb (Eds.), *Handbook of Social Capital*. Oxford: Oxford University Press.

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

[47] Ensmenger, N. (2010). *The Computer Boys Take Over: Computers, Programmers, and the Politics of Technical Expertise*. Cambridge, MA: MIT Press.

[48] Espinosa, M. and C. T. Stanton (2024). Training, communications patterns, and spillovers inside organizations. Revise & Resubmit, *Journal of Political Economy*.

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

[50] Gibbons, R. and J. Roberts (2014). Economic theories of incentives in organizations. In R. Gibbons and J. Roberts (Eds.), *The Handbook of Organizational Economics*. Princeton, NJ: Princeton University Press.

[51] Gorman, E. H. (2005). Gender stereotypes, same-gender preferences, and organizational variation in the hiring of women: Evidence from law firms. *American Sociological Review 70*(4), 702–728.

[52] Graham, P. (2004). *Hackers & Painters: Big Ideas from the Computer Age*. Sebastopol, CA: O'Reilly Media.

[53] Greenwald, A. G. and M. R. Banaji (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review 102*(1), 4–27.

[54] Grusky, D. B. (2005). Foundations of a neo-durkheimian class analysis. In E. O. Wright (Ed.), *Approaches to Class Analysis*, pp. 51–81. Cambridge University Press.

[55] Highhouse, S. (2008). Stubborn Reliance on Intuition and Subjectivity in Employee Selection. *Industrial and Organizational Psychology: Perspectives on Science and Practice 1*(3), 333–342.

[56] Hinds, P., K. M. Carley, D. Krackhardt, and D. Wholey (2000, February). Choosing work-group members: Balancing similarity, competence, and familiarity. *Organizational Behavior and Human Decision Processes 81*(2), 226–251.

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

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

[59] 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.

[60] Jodhka, S. S. (2015). *Caste in Contemporary India*. Routledge India.

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

[62] Kahn, C. and G. Huberman (1988). Two-sided uncertainty and "up-or-out" contracts. *Journal of Labor Economics 6*(4), 423–444.

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

[64] Kang, S. K., K. A. DeCelles, A. Tilcsik, and S. Jun (2016). Whitened resumes: Race and self-presentation in the labor market. *Administrative Science Quarterly 61*(3), 469–502.

[65] Kashyap, R. and F. Verkroost (2021). Analysing global professional gender gaps using LinkedIn advertising data. *EPJ Data Science 10*(39), 1–32.

[66] Khan, S. (2012). The sociology of elites. *Annual Review of Sociology 38*, 361–377.

[67] Khan, S. R. (2011). *Privilege: The Making of an Adolescent Elite at St. Paul's School*. Princeton, NJ: Princeton University Press.

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

[69] Kraus, M. W., J. W. Park, and J. J. Tan (2017). Signs of social class: The experience of economic inequality in everyday life. *Perspectives on Psychological Science 14*(3), 422–435.

[70] Kuncel, N. R. and S. A. Hezlett (2007). Standardized tests predict graduate students' success. *Science 315*(5815), 1080–1081.

[71] Kuncel, N. R., D. M. Klieger, B. S. Connelly, and D. S. Ones (2013). Mechanical versus clinical data combination in selection and admissions decisions: A meta-analysis. *Journal of Applied Psychology 98*(6), 1060–1072.

[72] Kuncel, N. R., R. J. Kochevar, and D. S. Ones (2014). A meta-analysis of letters of recommendation in college and graduate admissions: Reasons for hope. *International Journal of Selection and Assessment 22*(1), 101–107.

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

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

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

[76] Lareau, A. (2003). *Unequal Childhoods: Class, Race, and Family Life*. Berkeley, CA: University of California Press.

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

[78] Lazear, E. P. and S. Rosen (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy 89*(5), 841–864.

[79] Lazear, E. P., K. L. Shaw, and C. T. Stanton (2015). The value of bosses. *Journal of Labor Economics 33*(4), 823–861.

[80] Lenski, G. E. (1954). Status crystallization: A non-vertical dimension of social status. *American Sociological Review 19*(4), 405–413.

[81] Levashina, J., C. J. Hartwell, F. P. Morgeson, and M. A. Campion (2014). The structured employment interview: Narrative and quantitative review of the research literature. *Personnel Psychology 67*(1), 241–293.

[82] Levon, E., D. Sharma, D. J. L. Watt, A. Cardoso, and Y. Ye (2021). Accent bias and perceptions of professional competence in England. *Journal of English Linguistics 49*(4), 355–388.

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

[84] Mamidi, P. (2011). Strategic manipulation of caste identity signals. Doctoral Dissertation, Oxford University.

[85] Martellini, P., T. Schoellman, and J. Sockin (2024, February). The global distribution of college graduate quality. *Journal of Political Economy 132*(2), 434–483.

[86] Meehl, P. E. (1954). *Clinical versus Statistical Prediction: A Theoretical Analysis and a Review of the Evidence*. University of Minnesota Press.

[87] Mero, N. P. and S. J. Motowidlo (1995). The effects of rater accountability on the accuracy and impression-management bias of performance appraisal. *Journal of Applied Psychology 80*(4), 517–524.

[88] Michelman, V., J. Price, and S. D. Zimmerman (2022). Old boys' clubs and upward mobility among the educational elite. *The Quarterly Journal of Economics 137*(2), 845–909.

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

[90] Minni, V. (2024b). Making the invisible hand visible: Managers and the allocation of workers to jobs. R&R, *Quarterly Journal of Economics*.

[91] Minni, V., K.-T. Nguyen, and H. Sarsons (2025). Managers and the cultural transmission of gender norms. Working Paper.

[92] Murphy, S. C., D. M. Klieger, M. J. Borneman, and N. R. Kuncel (2009). The predictive power of personal statements in admissions: A meta-analysis and cautionary tale. *Educational and Psychological Measurement 69*(6), 967–982.

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

[94] Page, S. E. (2007). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton, NJ: Princeton University Press.

[95] Parmar, R. (2020). Transacting caste in modern times: Changing social identity through surnames in urban Gujarat. *Contemporary Voice of Dalit 12*(2), 220–233.

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

[97] Puri, S. (2025, February). Capitalism with Indian characteristics. Stigler Center Lecture Series, University of Chicago Booth School of Business. Accessed: 2025-02-11.

[98] Quillian, L., J. J Lee, and M. Oliver (2020). Evidence from field experiments in hiring shows substantial additional racial discrimination after the callback. *Social Forces 99*(2), 732–759.

[99] Reagans, R. and E. W. Zuckerman (2001, July–August). Networks, diversity, and productivity: The social capital of corporate R&D teams. *Organization Science 12*(4), 502–517.

[100] Reskin, B. (2012). The race discrimination system. *Annual Review of Sociology 38*, 17–35.

[101] Ridgeway, C. L. (2019). *Status: Why Is It Everywhere? Why Does It Matter?* New York, NY: Russell Sage Foundation.

[102] Rivera, L. A. (2011). Ivies, extracurriculars, and exclusion: Elite employers' use of educational credentials. *Research in Social Stratification and Mobility 29*(1), 71–90.

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

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

[105] Rivera, L. A. and A. Tilcsik (2016). Class advantage, commitment penalty: The gendered effect of social class signals in an elite labor market. *American Sociological Review 81*(6), 1097–1131.

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

[107] 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.

[108] Sadun, R. (2023). CEOs and firm performance. *NBER Reporter* (1), 12–16.

[109] Sadun, R., N. Bloom, and J. Van Reenen (2017, September-October). Why do we undervalue competent management? Neither great leadership nor brilliant strategy matters without operational excellence. *Harvard Business Review 95*(5), 120–127. Winner of 59th Annual HBR McKinsey Award.

[110] Sadun, R., J. A. Tamayo, J. Ramirez, B. S. Diaz, and A. Neyra-Nazarrett (2024, April). Training within firms. Working Paper 33670, National Bureau of Economic Research.

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

[112] Sarsons, H. (2023). Interpreting Signals in the Labor Market: Evidence from Medical Referrals. R&R, *Review of Economic Studies*.

[113] 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.

[114] Savage, M. (2000). *Class Analysis and Social Transformation*. Sociology and Social Change. Buckingham: Open University Press.

[115] Savage, M. (2015). *Social Class in the 21st Century*. Pelican Books.

[116] Savage, M., F. Devine, N. Cunningham, M. Taylor, Y. Li, J. Hjellbrekke, B. Le Roux, S. Friedman, and A. Miles (2013). A new model of social class? Findings from the bbc's great british class survey experiment. *Sociology 47*(2), 219–250.

[117] Savage, M., F. Devine, N. Cunningham, M. Taylor, Y. Li, J. Hjellbrekke, B. Le Roux, S. Friedman, and A. Miles (2015). On social class, anno 2014. *Sociology 49*(6), 1011–1030.

[118] Stanovich, K. E. and R. F. West (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences 23*(5), 645–665.

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

[120] Stansbury, A. and K. Rodriguez (2024b, December). The class gap in career progression: Evidence from US academia. Presented at the NBER Organizational Economics Working Group, Fall 2024.

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

[122] Tetlock, P. E. (2005). *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton, NJ: Princeton University Press.

[123] Tetlock, P. E. and B. A. Mellers (2011). Structuring accountability systems in organizations: Key trade-offs and critical unknowns. In B. Fischhoff and C. Chauvin (Eds.), *Intelligence Analysis: Behavioral and Social Scientific Foundations*, Chapter 11, pp. 249–270. Washington, DC: National Academies Press.

[124] Thorat, S. and P. Attewell (2007, October). The legacy of social exclusion: A correspondence study of job discrimination in India. *Economic and Political Weekly 42*(41), 4141–4145.

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

[126] Turco, C. J. (2010). Cultural foundations of tokenism: Evidence from the leveraged buyout industry. *American Sociological Review 75*(6), 894–913.

[127] Turkle, S. (1984). *The Second Self: Computers and the Human Spirit*. New York: Simon and Schuster.

[128] Weber, M. (1922). *Economy and Society: An Outline of Interpretive Sociology*, Volume 1. Berkeley, CA: University of California Press.

[129] Wilkerson, I. (2020). *Caste: The Origins of Our Discontents*. New York: Random House.

[130] Wilmers, N., S. Wang, and L. Zhang (2025). Remote work and careers. SSRN working paper.

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

[132] Zimmerman, S. D. (2019). Elite colleges and upward mobility to top jobs and top incomes. *American Economic Review 109*(1), 1–47.

ONLINE APPENDIX

A Caste as Social Status Under Affirmative Action

India's caste system provides a uniquely precise measure of social status, offering clear administrative categories for analyzing class-based inequalities. Affirmative action policies formally recognize caste as a comprehensive marker of historical advantage and disadvantage, codifying how caste membership systematically shapes access to economic resources, education, social networks, and elite employment opportunities—dimensions that align with Bourdieu's (1986) conceptualization of social class. This administrative clarity enables identification of status-based discrimination mechanisms that remain difficult to isolate in class analyses in other contexts.

Jati, Varna, and Modern Caste Classifications. India's modern caste system originates from historical occupational hierarchies. At its most detailed level, the system comprises thousands of hereditary occupational groups known as *Jatis*, historically linked to trades such as merchants, priests, warriors, and laborers. These *Jatis* are grouped into broader categories called *Varnas*: Brahmins (priests), Kshatriyas (warriors), Vaishyas (merchants), Shudras (laborers), plus a historically excluded group called "untouchables" (Dalits). The government maintains official lists further aggregating and classifying disadvantaged communities as "Scheduled Castes" (SCs), "Scheduled Tribes" (STs), and "Other Backward Classes" (OBCs)—groups defined as socially and educationally disadvantaged (Puri, 2025).

This paper adopts two broad categories aligned with India's affirmative action framework: advantaged ("upper") and disadvantaged ("lower") castes. By disadvantaged castes, I refer explicitly to SCs, STs, and OBCs. This administrative categorization recognizes caste as fundamentally intertwined with economic marginalization and social status, paralleling Bourdieusian concepts of class.

Historical Development of Caste-based Affirmative Action Policies. Affirmative action for disadvantaged communities emerged during British colonial rule, beginning with the Government of India Act of 1919's reserved legislative seats for "depressed classes," defined as those who were economically and socially underprivileged. The 1935 Government of India Act officially designated these groups as Scheduled Castes. After independence from British rule was achieved in 1947, India's Constitution institutionalized quota-based affirmative action policies in legislatures, public higher education, and government employment for Scheduled Castes and Scheduled Tribes. Following the Mandal Commission's recommendations in the 1980s, quotas expanded to additionally include Other Backward Classes, raising reservation levels to nearly 50% for all disadvantaged communities combined (Bayly, 2008).

Caste as Proxy for Social Status. Although caste originated as a hereditary hierarchy, Indian policymakers now define caste categories as comprehensive indicators of socioeconomic "backwardness"—encompassing economic, educational, and social disadvantages. This approach recognizes caste as encapsulating multiple dimensions of social status: financial resources, educational credentials, social networks, and cultural capital. The Mandal Commission

expanded caste categories beyond hereditary lineage to include socioeconomic indicators when defining "other" backward *classes* (i.e., OBCs), further establishing caste as a consolidated marker of social status.

Limited Scope of Affirmative Action in the Private Sector. India's affirmative action policies do not extend to private-sector employment, including the multinational corporations studied here. Private-sector roles typically offer higher wages and career growth but disproportionately exclude disadvantaged castes (Subramanian, 2019). This parallels contexts such as the U.S., where firms emphasize racial and gender diversity but rarely address socioeconomic or class diversity (Stansbury and Rodriguez, 2024a,b).

B Representativeness of the Job Fair at the Elite College

The studied on-campus job fair provides a highly representative window into how India's elite-college graduates transition into elite private-sector jobs, supported by four key observations:

1. **Near-identical fairs across elite institutions.** India's premier technical institutions, established in the 1950s, consciously modeled their practices after prominent Western universities such as MIT, Stanford, and Harvard (Datta, 2017), including their placement office structures that had been developing in the U.S. since the 1930s (Anderson, 1938). Subsequently, other elite Indian institutions largely replicated these early pioneers' models, creating a cascading standardization effect. Consequently, these institutions now share strikingly similar academic calendars, faculty-to-student ratios, curricula, and, critically, standardized placement office procedures. Today, job fairs across leading Indian universities feature virtually identical recruitment stages, timelines, and employer pools (see, e.g., Central Placement Cell, Delhi University).

2. **Graduates predominantly flow into elite private-sector jobs.** Graduates from India's elite colleges overwhelmingly pursue high-paying entry-level roles in multinational corporations and prominent Indian firms. Public-sector employment remains a clear secondary choice, reflecting a long-standing preference among top-tier college graduates for careers in consulting, technology, finance, and management (Subramanian, 2019).

3. **Elite firms recruit almost exclusively from elite campuses.** The multinational firms and prestigious Indian companies studied here systematically target the same elite Indian institutions year after year. Collectively, these recruiters visit essentially the same set of top-tier colleges each year, with each institution typically enrolling more than 1,000 students. This concentrated recruiting strategy ensures a highly consistent candidate pool across cohorts. Prior studies further corroborate that elite firms' hiring remains strongly concentrated within these elite campuses (Jodhka and Naudet, 2019).

4. **Minimal student withdrawal from the job fair.** Participation in the on-campus job fair is near-universal among graduating students. Less than 5 percent of final-year students voluntarily deregister from the formal placement process, a proportion consistent across similar elite institutions (e.g., Mamgain, 2019). Such low withdrawal rates underscore the central role of these standardized recruitment fairs in determining employment outcomes for India's most academically competitive graduates.

Collectively, these four facts demonstrate that the focal job fair closely mirrors elite hiring pathways across India's top universities.

C Timeline and Screening Rounds

Every job in the sample follows four sequential screening stages, eliminating candidates at each step. The following list describes the screening rounds and the timeline:

1. **Application screening (September).** Recruiters conduct an initial screening, primarily based on GPA, major, degree, and standout résumé items (e.g., prestigious internships, performance at inter-college competitions).

2. **Written aptitude tests (September–November).** Online technical assessments vary by role: usually coding exercises for tech jobs, and case-based tests for quantitative consulting positions. Firms advertising multiple positions usually administer one common test (e.g., Samsung may use single exam for all advertised positions).

3. **Group debates (September–November).** A non-technical evaluation of soft skills:

   (a) Groups of 30-35 students, split into two teams, debate a general prompt (e.g., "Is social media a boon or a bane?" or "What are the main bottlenecks faced by Indian tech startups?"), while recruiters observe communication, teamwork, and confidence.

   (b) Firms often conduct a single group debate for all advertised roles—for example, Google may simultaneously assess applicants for a program manager and software engineering position. This common screening approach suggests that firms value similar socio-emotional competencies—communication, teamwork, confidence—across diverse technical roles.

4. **Personal interviews (December–January).** Also called HR interviews, these are subjective assessments centering on candidate "fit":

   (a) Interviews occur simultaneously in multiple parallel batches, each handled by different members of the recruiting team.

   (b) As with earlier rounds, a single interview panel commonly covers multiple job titles (e.g., Uber may interview data scientist and front-end engineering candidates via the same interview panel).

   (c) Because almost all roles are India-based, interviewers are almost always Indian hiring managers (Section 3.4).

D Benchmarking Employment-Platform Coverage: External Evidence

Large-scale validations suggest LinkedIn already reaches roughly 80% of U.S. bachelor's-degree holders (Wilmers et al., 2025). Coverage is markedly higher in technical or STEM fields. LinkedIn's own *Pathways into STEM Employment* study reports that 92.7% of 2022 STEM

master's graduates (virtually all STEM users with active profiles) interacted with at least one STEM job on the site, a gain of nearly 13 percentage points over the U.S. bachelor-level average of 80% (Baird et al., 2023). Cross-country evidence reveals this STEM overrepresentation is not unique to the United States. In their 92-nation ad-audience study, Kashyap and Verkroost (2021) show that for India the *Information and Communication* sector accounts for 32% of female and 27% of male LinkedIn users but only about 1% of the professional labor force logged by the International Labour Organization, an enrichment of more than 25 percentage points that mirrors the global STEM premium they document.

Industry-specific data reinforce this pattern: Tambe et al. (2021) find IT employees constitute 5% of LinkedIn's U.S. workforce versus 2% of total U.S. employment in Bureau of Labor Statistics data, meaning IT workers are 2.5 times overrepresented on the platform. Given that about 35% of U.S. adults use LinkedIn overall (Gottfried, 2024), this 2.5x overrepresentation implies approximately 90% of IT professionals are on LinkedIn ($35\% \times 2.5 \approx 90\%$), approaching near-census coverage. This calculation assumes the 35% overall-population adoption rate applies uniformly across occupations; in reality LinkedIn usage is higher among college-educated STEM professionals (LinkedIn Economic Graph, 2018), so 90% is likely a conservative estimate.

Complementing the evidence above, elite STEM-focused universities also push LinkedIn coverage to near-census levels: Stanford's official count lists more than 220,000 living alumni, yet LinkedIn shows about 314,000 profiles, which after de-duplication and mis-tag removal using standard methods yields a match rate exceeding 95%.[1] MIT demonstrates similar coverage: while its alumni office reports around 149,000 living graduates, LinkedIn captures roughly 233,000 MIT profiles—again indicating over 95% real-match coverage.[2]

Taken together, these findings establish LinkedIn as a near-comprehensive sampling frame for studying highly educated STEM professionals. Given that our college's STEM talent ranks on par with that of the world's leading institutions (Martellini et al., 2024), a well over 90% in-sample match rate at our campus is consistent with these external benchmarks. Indeed, the third-party matched dataset achieves a nearly 98% match rate for graduates (higher by about 9% than other benchmarks that hover around 90%), aided by systematic institutional identifiers that follow predictable patterns within each major-degree-year, which greatly boosts match precision. This near-complete coverage effectively captures the full universe of our STEM graduates' career trajectories.

E Robustness to Omitted Variable Bias: Oster Bounds

E.1 Motivation and Methodology

A central concern in our analysis of caste-based discrimination is the potential for omitted variable bias. While we control for a rich set of observable characteristics including individual qualifications, job characteristics, and cohort fixed effects, unobservable factors correlated with both caste and hiring outcomes could bias our estimates. To address this concern, we implement the bias-correction methodology developed by Oster (2019), building on the insights of Altonji

[1] See Zhang and Vucetic (2016) for methods to de-duplicate such data.

[2] Stanford Alumni Association; LinkedIn People (Stanford University); MIT Alumni Association; LinkedIn People (MIT). Even outside STEM fields, LinkedIn achieves strong coverage: MBA follow-ups routinely match 85–90% of graduates from many cohorts to profiles (Hampole et al., 2024).

et al. (2005).

The Oster method provides a systematic approach to bounding treatment effects in the presence of unobservables. The key insight is that by examining how coefficient estimates and R-squared values change as controls are added, we can infer the likely bias from omitted variables under reasonable assumptions about selection on observables versus unobservables. Specifically, the method requires two key parameters:

1. δ **(delta):** The relative degree of selection on unobservables versus observables.
2. $R^*_{\mathbf{max}}$**:** The R-squared from a hypothetical regression including all relevant controls (both observable and unobservable).

E.2 Parameter Choices

Following Oster (2019), we adopt conservative parameter values:

1. **Choice of $\delta = 1$:** We assume that selection on unobservables is as important as selection on observables. While δ could theoretically exceed 1 (implying unobservables matter more than observables), this seems implausible in our context. We already control for a rich set of qualifications (e.g., pre-college test scores, entrance exam score ranks, college GPA, prior labor market experience) used in hiring decisions. Any remaining unobserved characteristics are unlikely to be more predictive of hiring outcomes than these direct measures of candidate quality. Moreover, some outcomes (e.g., promotions) approach an upper bound (≈ 0.70), suggesting a "boundary bias," so even large increases in δ have limited leverage to alter the estimated gap.
2. **Choice of $R^*_{\mathbf{max}} = 1.3\tilde{R}$:** Following Oster's recommendation, we set the *maximum R-squared* to 1.3 times the R-squared from our controlled regression ($\tilde{R}$). This assumes that including all unobservables would increase explanatory power by 30%, which is substantial given our already comprehensive set of controls. As Oster (2019) notes, this is a conservative benchmark that has been validated across numerous empirical applications.

Parameter choices are conservative. The conservatism of these choices is particularly relevant in the discrimination context. As demonstrated in applications like Stansbury and Rodriguez (2024a), when studying disadvantaged groups in elite settings, the observable measures of productivity and qualifications are often the very metrics used in selection processes. Thus, the scope for unobservables to explain substantial additional variation—and to be correlated with group membership in ways that would eliminate group disparities—is limited.

E.3 Results

Table A presents our main caste coefficients presented in Tables 6 and 7 alongside their Oster-adjusted counterparts. The stability of our coefficients under the Oster correction implies that our main findings are not artifacts of omitted variable bias. For the critical interview success outcome, the coefficient moves only marginally from -0.294 to -0.296. Similarly, the promotion rate gaps remain virtually unchanged, with most coefficients shifting by less than 0.01. Moreover, for promotion regressions, the bounded nature of the outcome (promotions cannot exceed

100%) creates boundary bias that inherently limits the potential for unobservables to explain additional variation.

Indeed, in some specifications (notably test success), the Oster-adjusted coefficient actually becomes more positive, suggesting that if anything, unobservables may be masking rather than creating the disparities we observe. This pattern is consistent with disadvantaged caste members potentially having unobserved positive characteristics (such as greater perseverance or motivation required to reach elite positions) that partially offset discrimination.

Table A: **Caste Estimates: OLS vs. Oster-Adjusted**

	Disadvantaged Caste	
	OLS	Oster-Adjusted
	(1)	(2)
Panel A: Success Rates (Table 6)		
Application Success	-0.0025*	0.009
	(0.0014)	(0.0013)
Test Success	-0.0002	0.010***
	(0.0018)	(0.002)
Debate Success	-0.0005	-0.002
	(0.0027)	(0.003)
Interview Success	-0.2938***	-0.296***
	(0.0048)	(0.005)
Panel B: Promotion Rates (Table 7)		
Promoted (All)	0.137***	0.142***
	(0.025)	(0.026)
Promoted (w/ controls)	0.096***	0.093***
	(0.026)	(0.028)
Promoted at Top 75%	0.104***	0.108***
	(0.029)	(0.031)
Promoted at Top 50%	0.112***	0.122***
	(0.038)	(0.037)
Promoted at Top 40%	0.136***	0.141***
	(0.043)	(0.041)
Promoted at Top 30%	0.147***	0.157**
	(0.054)	(0.058)
Promoted at Top 20%	0.190**	0.196**
	(0.073)	(0.082)
Promoted at Top 15%	0.204**	0.206**
	(0.084)	(0.082)
Promoted at Top 10%	0.084	0.074
	(0.109)	(0.131)
Promoted (Client Facing)	0.107**	0.106***
	(0.044)	(0.039)
Promoted (Non-Client Facing)	0.082**	0.075**
	(0.032)	(0.036)

Notes: Column (1) reports OLS estimates with full controls. Column (2) reports bias-adjusted estimates using the Oster (2019) methodology with $\delta = 1$ and $R^*_{\max} = 1.3\tilde{R}$, where $\tilde{R}$ denotes the R-squared from the controlled regression. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. All specifications include cohort fixed effects, job fixed effects, and other controls as specified in the main tables.

F Qualitative evidence supporting caste revelation at interviews

Qualitative evidence further clarifies why caste becomes most legible to recruiters during personal interviews. As discussed quantitatively in Section 4.6.1, surnames alone rarely disclose caste in urban Indian settings; instead, status cues often emerge through richer *social interactions*. Since written tests are administered electronically, recruiters have little information beyond test scores and the résumés they already possess. Group debates provide limited scope for caste inference: candidates debate standardized prompts in large groups with minimal individualized questioning, and neither physical appearance nor linguistic cues reliably indicate caste. Indeed, anthropological research consistently finds phenotypic traits like skin color are unreliable caste indicators among urban, educated populations (Sebring, 1969; Mishra, 2015; Parameswaran and Cardoza, 2015). The human-rights organization *Human Rights Watch* concurs, stating that "Lower-caste communities are almost invariably indistinguishable in physical appearance from higher-caste communities" (Human Rights Watch, 2001, p. 2). Similarly, varieties of Indian English spoken by the college-educated typically reflect regional or educational backgrounds rather than caste, as English-medium schooling has expanded significantly beyond advantaged castes due to decades of affirmative action (Wiltshire, 2020; Sharma, 2017; Highet, 2023). Thus, group debates offer employers minimal caste-identifying information.

In contrast, personal interviews substantially shift the evaluative context. They not only *amplify* earlier signals such as schooling pedigree, résumé details, and debate performances through deeper discussions, but also reveal *new* information previously unavailable to recruiters. Interviews elicit richer background information through individualized conversations about personal interests, leisure activities, extracurricular involvement, and peer networks (e.g., Rivera, 2012, 2015), which are strongly correlated with caste-linked social and cultural capital (Bourdieu, 1986), such as networks, resources, and cultural tastes. Such characteristics are well-documented caste correlates in elite Indian contexts (cf. Jodhka and Newman, 2007; Deshpande, 2011; Jodhka, 2015). Additionally, caste identities frequently become salient through social interactions in organizational settings, as managers learn subtle cues from conversations about personal background, residential origins, familial occupations, or shared acquaintances (Bapuji et al., 2023). Integrating these nuanced signals with earlier observations enables recruiters to more accurately infer caste or social class, aligning precisely with the timing and intensity of caste penalties documented quantitatively (Section 4.6).

G Extending the Simple Model of Late-Stage Discrimination with Imperfect Information

In Section 4.7.1, we assumed that employers observe each candidate's true screening-round performance perfectly. Here, I show that *all* key insights—particularly how variation in job selectivity identifies discrimination versus ability deficits—remain valid even when employers observe *noisy* performance signals.

Setup. As before, there are two groups (castes) $g \in \{0, 1\}$, where $g = 1$ denotes disadvantaged castes and $g = 0$ denotes advantaged castes, a continuum of jobs $j \in \mathcal{J}$, and 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

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

where $M_{s,i,g}$ is the score of candidate i who belongs to group g at round s, $\mathcal{I}_{s-1}$ denotes the employer's *information set* after observing ability signals in rounds $1, \ldots, s-1$, and $\tau_{s,j}$ is the job-specific threshold for round s. As before, let $\tau_{4,j}^{(g)}$ be the *group-specific* interview threshold at job j. If caste becomes salient at the interview for group 1, a higher threshold ($\tau_{4,j}^{(1)} > \tau_{4,j}^{(0)}$) indicates discrimination. 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 satisfying:

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

(ii) $\eta_{s,i,1} \overset{d}{=} \eta_{s,i,0}$ for every round s, so that for any subset of candidates (including those who have passed earlier rounds) there is no systematic caste-level difference in noise distributions. Thus, while noise might correlate across rounds or candidates (e.g., candidate performance may be interdependent across stages), it never introduces *systematic* group-level biases. Under these conditions, persistent group differences cannot result from noise alone.

Assumptions. Throughout this appendix we maintain the assumptions laid out in Section 4.7.1 pertaining to regularity of the screening score distributions (Assumption 2), monotonicity of screening thresholds (Assumption 3), common talent pool across jobs (Assumption 4) and the screening score distribution having light tails (Assumption 5). Given the common pool assumption across jobs (Assumption 4), each job draws candidates whose screening scores come from the same distributions $F_{s,g}$. Screening noise $\eta_{s,i,g}$ is therefore orthogonal to cross-job variation in $\tau_{s,j}$. Since screening score distributions are light-tailed (Assumption 5), each $F_{s,g}$ has a diverging hazard rate ($h_{F_{s,g}}(\tau) \to \infty$). Assuming $\eta_{s,i,g}$ is light-tailed as well, each signal $M_{s,i,g} = A_{i,g} + \eta_{s,i,g}$ inherits that property, ensuring $\mathbb{E}[M - \tau \mid M \geq \tau] \to 0$, as shown in Online Appendix Section H.

Proposition 3 (No Early Gap Implies No Right-Tail Deficit Before Interviews). Consider any pre-interview stage $s < 4$. Under Assumptions 2–5, if disadvantaged castes (group $= 1$) truly have fewer high-ability candidates in the right tail, then for sufficiently high $\tau_{s,j}$,

$$P(\text{pass}_{s,j} \mid g = 1) - P(\text{pass}_{s,j} \mid g = 0) < 0.$$

Therefore, empirically observing no such gap—even among highly selective jobs—rules out a meaningful right-tail deficit for group 1 before interviews.

Sketch of Proof. Given a common talent pool across jobs (Assumption 4), all jobs sample from the same $F_{s,g}$. With noisy signals, passing round s depends on the posterior mean $\mathbb{E}[M_{s,i,g} \mid \mathcal{I}_{s-1}]$. For high $\tau_{s,j}$, employers require a sufficiently high $\mathbb{E}[M_{s,i,g} \mid \mathcal{I}_{s-1}]$. If group 1 has fewer

high-ability individuals, then because performance distributions are light-tailed (Assumption 5), the pass rate for group 1 falls *faster* than group 0 as $\tau_{s,j}$ rises. This mirrors the perfect-information argument in Proposition 1. Because noise is symmetric across groups, it cannot conceal large latent ability gaps at high cutoffs. Thus, the absence of early-stage gaps even at very selective jobs implies no meaningful right-tail deficit for disadvantaged castes. □

Given Proposition 3, any observed late-stage gap must stem either from newly revealed ability signals during interviews or from discrimination. To distinguish between these explanations under imperfect information, I examine how post-hire promotion patterns vary with job selectivity. Post-hire performance Y_i depends on latent ability $A_{i,g}$. Promotion requires $Y_i \geq \kappa_j$. Final-round passing demands $\mathbb{E}[M_{4,i,g} \mid \mathcal{I}_3] \geq \tau_{4,j}^{(g)}$.

Proposition 4 (Stable Promotion Gaps Identify Discrimination). Under Assumptions 2–5:

1. *No discrimination.* If $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)}$, then promotion gaps vanish as selectivity increases.

2. *Discrimination.* If $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} + \delta$ with $\delta > 0$, then group 1's average performance remains approximately δ higher, even in highly selective jobs.

Hence, persistent promotion gaps at high selectivity reject an unobserved-ability explanation in favor of discrimination.

Sketch of Proof. The argument parallels the one made in Proposition 2. By Proposition 5, the mean-residual life of both ability and noise vanishes at high cutoffs. Under imperfect information, stage-4 passage requires $\mathbb{E}[M_{4,i,g} \mid \mathcal{I}_3] \geq \tau_{4,j}^{(g)}$. If $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} + \delta$, then as job selectivity increases, *average ability of hires converges to that of the marginal hire* (hiring cutoff) and therefore the typical group 1 hires satisfy:

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

Because noise is identically distributed across groups, raising group 1's threshold by δ selects candidates with ability approximately δ higher even as job selectivity increases. If hiring thresholds were identical, performance gaps would have vanished with increasing job selectivity. □

H Derivation of Expected Performance Above Threshold

This appendix establishes the limit result used in Proposition 2,

$$\lim_{\tau \to \infty} \big(\mathbb{E}[M \mid M \geq \tau] - \tau \big) = 0.$$

The proof is self-contained and relies on a single, transparent tail assumption.

Setup. Let M be a real-valued random variable with CDF F, PDF f, survival function $S(a) = 1 - F(a)$, and hazard rate $h(a) = f(a)/S(a)$ (defined wherever $S(a) > 0$).

Assumption A. M has a continuously differentiable density, finite mean, and its hazard rate diverges:

$$\lim_{a\to\infty} h(a) = \infty.$$

This covers all light-tailed, log-concave distributions commonly used in practice (e.g. Normal, Log-Normal, Logit, Weibull with shape parameter > 1).

Proposition 5 (Mean of a Right-Truncated Distribution)**.** For any τ with $S(\tau) > 0$, define the mean residual life

$$m(\tau) \;\equiv\; \mathbb{E}[M - \tau \mid M \geq \tau] \;=\; \frac{\displaystyle\int_\tau^\infty S(u)\,du}{S(\tau)}.$$

Under Assumption A,

$$\lim_{\tau\to\infty} m(\tau) = 0, \quad \text{hence} \quad \mathbb{E}[M \mid M \geq \tau] = \tau + m(\tau) \;\to\; \tau.$$

Proof. Let

$$A(\tau) = \int_\tau^\infty S(u)\,du,$$

so $m(\tau) = A(\tau)/S(\tau)$. Then

$$A'(\tau) = -S(\tau), \qquad S'(\tau) = -f(\tau).$$

Differentiating $m(\tau) = A(\tau)/S(\tau)$ gives

$$m'(\tau) = \frac{A'(\tau)\,S(\tau) \;-\; A(\tau)\,S'(\tau)}{S(\tau)^2} = \frac{-S(\tau)^2 \;+\; f(\tau)\,A(\tau)}{S(\tau)^2} = -1 + \underbrace{\frac{f(\tau)}{S(\tau)}}_{h(\tau)}\,\underbrace{\frac{A(\tau)}{S(\tau)}}_{m(\tau)} = -1 + h(\tau)\,m(\tau).$$

Rearrange to

$$m'(\tau) + 1 \;=\; h(\tau)\,m(\tau) \quad \to \quad m(\tau) = \frac{1 + m'(\tau)}{h(\tau)}.$$

Since $m(\tau) \geq 0$ and m is non-increasing, $m'(\tau) \leq 0$, so

$$0 \;\leq\; 1 + m'(\tau) \;\leq\; 1.$$

Under Assumption A, $h(\tau) \to \infty$, hence

$$0 \;\leq\; m(\tau) = \frac{1 + m'(\tau)}{h(\tau)} \;\to\; 0,$$

which completes the proof. □

Discussion. The key requirement is that the hazard rate diverges. Heavy-tailed distributions with $h(\tau) \to 0$ (e.g., Pareto) fail this property and yield $m(\tau) \to \infty$ instead. In our context,

where performance scores are bounded and therefore light-tailed, the limit result holds, justifying the asymptotic arguments in Proposition 2.

I Cultural Fit in Final-Round Interviews at Elite Multinationals

Final-round interviews at elite multinationals often hinge on interviewers' discretionary judgments of *cultural fit*. In her seminal work on hiring practices at elite firms, Rivera (2012) shows that employers practice a form of "cultural matching," favouring candidates whose background, tastes, and self-presentation feel similar to theirs. In her paper, she surveys U.S. professional-service firms, where more than half of interviewers ranked "cultural fit" as *the single most important* criterion in the final round interviews, ahead of analytical skill or academic pedigree. A law-firm partner captured the prevailing logic:

> "In our new associates we are, first and foremost, looking for *cultural* compatibility. Someone who will . . . *fit in*." (Rivera, 2012, p. 1006)

"Would I Want to Be Stuck in an Airport With Them?"

Rivera (2012) found that interviewers often invoke the *airport test*: would I enjoy being stranded with this candidate for hours? A senior banker explained:

> "A lot of this job is attitude, not aptitude . . . fit is really important. So you can be the smartest guy ever, but I don't care. I need to be comfortable working everyday with you, then getting stuck in an airport with you, and then going for a beer after. You need chemistry. Not only that the person is smart, but that you *like* him." (Rivera, 2012, p. 1008)

Such remarks show how late-stage evaluation often becomes an assessment of social rapport and who "belongs" rather than pure competence (Lamont, 2009). To locate common ground quickly, interviewers scour the résumé for shared experiences—college sports, common extracurricular, personal interests, study-abroad destinations, niche hobbies—and open the interview with informal chat. If a spark appears, enthusiasm often translates into advocacy in the hiring committee. Conversely, when overlap is missing, evaluators may withdraw support even when the candidate excels technically (Rivera, 2012, pp. 1013–1019).

Hiring as Cultural Matching

The traits that signal "fit" are deeply entwined with social class and cultural capital (Bourdieu, 1986). Playing squash or lacrosse can mark a candidate as privileged in one firm and "too snooty" in another, depending on the organization's self-image. Rivera documents a consulting director who excluded a technically strong candidate because they were "more rough and tumble" (Rivera, 2012, p. 1009). Similar dynamics appear in other sectors: Turco's (2010) ethnography of a leveraged-buyout firm showed women struggling to meet an aggressive, sport-centred masculine ideal, while Erickson (1996) showed that shared leisure pursuits foster stronger workplace ties and, ultimately, advancement. Crucially, as Rivera (2012) demonstrates, professionalism and client-readiness are often evaluated *separately* from fit. As one consultant noted:

> "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." (Rivera, 2012, p. 1007)

Recruiters believed that while polish, presence, and presentation style could be taught, "fit" cannot. Because the finalist pool is already academically and technically elite, these subtle cultural cues decisively influence the offer decision (Lamont, 1992).

Implications for Caste

Our quantitative results show that caste penalties emerge almost exclusively at final round personal interviews. The qualitative literature helps explain why. If lower-caste candidates have had less exposure to elite cultural codes—speech patterns, leisure interests, informal networks—they may appear a weaker "fit'' despite equal credentials. Interviewers' reliance on self-referential judgments thus reinscribes social-status hierarchies under an ostensibly meritocratic system (Neckerman and Kirschenman, 1991; Deshpande, 2011). As Rivera (2012) concludes, employers often "constructed and assessed merit in their own image, believing that culturally similar applicants were better candidates'' (p. 1018).

Supplementary References

Altonji, J. G., T. E. Elder, and C. R. Taber (2005). Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *Journal of Political Economy 113*(1), 151–184.

Anderson, R. N. (1938). The work of the placement officer. *Teachers College Record 40*(1), 1–7.

Baird, M., N. Gahlawat, R. Hood, S. Lara, and P. Ko (2023). Pathways into STEM jobs. Technical report. LinkedIn Economic Graph, November 2023.

Bapuji, H., K. Gupta, S. Chrispal, and T. Roulet (2023). What managers everywhere must know about caste. *MIT Sloan Management Review Winter*. Accessed: 2024-04-22.

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

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.

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

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

Erickson, B. H. (1996). Culture, class, and connections. *American Journal of Sociology 102*(1), 217–251.

Gottfried, J. (2024, January). Americans' social media use. Technical report, Pew Research Center.

Hampole, M., F. Truffa, and A. Wong (2024). Peer effects and the gender gap in corporate leadership: Evidence from MBA students. R&R, *Quarterly Journal of Economics*.

Highet, K. (2023). "RIP English": Race, Class and 'Good English' in India. *Journal of Linguistic Anthropology 33*(2), 184–201.

Human Rights Watch (2001). Caste discrimination: A global concern. HRW Briefing Paper. See ch. 2, p. 2: "Lower-caste communities are almost invariably indistinguishable in physical appearance from higher-caste communities".

Jodhka, S. S. (2015). *Caste in Contemporary India*. Routledge India.

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

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.

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

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

LinkedIn Economic Graph (2018, August). Linkedin workforce report: United states. Technical report, LinkedIn Corporation.

Mamgain, R. P. (2019). *Formal Labour Market in Urban India: Job Search, Hiring Practices and Discrimination*. SAGE Publications Pvt. Ltd.

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

Neckerman, K. M. and J. Kirschenman (1991). Hiring strategies, racial bias, and inner-city workers. *Social Problems 38*(4), 433–447.

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

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.

Puri, S. (2025, February). Capitalism with Indian characteristics. Stigler Center Lecture Series, University of Chicago Booth School of Business. Accessed: 2025-02-11.

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

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

Sebring, J. M. (1969). Caste indicators and caste identification of strangers. *Human Organization 28*(3), 199–207.

Sharma, D. (2017). English in India. In A. Bergs and L. Brinton (Eds.), *The History of English: Volume 5-Varieties of English*, pp. 311–329. Boston: De Gruyter Mouton.

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

Stansbury, A. and K. Rodriguez (2024b, December). The class gap in career progression: Evidence from US academia. Presented at the NBER Organizational Economics Working Group, Fall 2024.

Subramanian, A. (2019). *The Caste Of Merit: Engineering Education in India*. Harvard University Press.

Tambe, P., L. Hitt, D. Rock, and E. Brynjolfsson (2021). Digital capital and superstar firms. R&R, *Management Science*.

Wilmers, N., S. Wang, and L. Zhang (2025). Remote work and careers. SSRN working paper.

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

Zhang, S. and S. Vucetic (2016). Sampling bias in Linkedin: A case study. In *Proceedings of the 25th International Conference Companion on World Wide Web*, pp. 145–146.

J Additional 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 (OLS)	-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: This table reports estimates from earnings regressions for graduates who obtained jobs through the college's job fair. The dependent variable is log(salary) of the accepted job. Controls include pre-college skills, GPA, prior labor market experience, major, degree, and other characteristics. Regressions include year fixed effects. Column (1) shows estimates without controls, Column (2) adds controls, and Columns (3)-(4) show separate estimates for client-facing and non-client-facing jobs with controls. Job fixed effects are omitted as salaries are identical within jobs regardless of caste. Bias-adjusted coefficients follow Oster (2019), setting $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$ (assuming unobservables match observables in selection intensity and allowing total explanatory power to exceed observed R^2 by 30%). For detailed discussion of these parameter choices, see Appendix E. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table OA.2: **Caste Differences in Mean Salary at Each Job Search Stage in the Full Sample**

	Dependent variable:					
	log (Avg. Salary)					
	App	Test	Debate	Interview	Offer	Choice
	(1)	(2)	(3)	(4)	(5)	(6)
Disadv. Caste (OLS)	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: This table reports earnings regressions for students still in contention at each job search stage. Each column represents a different stage: (1) application, (2) written test, and subsequent rounds. The dependent variable is log(average salary), where the average is taken over the posted salaries of all jobs a student remains eligible for at that stage (for the "Choice" column this collapses to the salary of the accepted job). Controls include pre-college skills, GPA, prior labor market experience, major, degree, and other characteristics. All regressions control for pre-college skills, GPA, prior labor market experience, major, and degree. Regressions include year fixed effects. Job fixed effects are omitted as salaries are identical within jobs regardless of caste. Bias-adjusted coefficients follow Oster (2019), setting $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$ (assuming unobservables match observables in selection intensity and allowing total explanatory power to exceed observed R^2 by 30%). For detailed discussion of these parameter choices, see Appendix E. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table OA.3: **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: This table reports caste differences in success rates at each hiring stage for client-facing jobs. The unit of observation is student-job-year, conditional on remaining in contention at each stage. All regressions include controls as well as job and year fixed effects. When firms conduct common screening rounds for multiple positions, each position is treated as a separate job. Bias-adjusted coefficients follow Oster (2019), setting $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$ (assuming unobservables match observables in selection intensity and allowing total explanatory power to exceed observed R^2 by 30%). For detailed discussion of these parameter choices, see Appendix E. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table OA.4: **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: This table reports caste differences in success rates at each hiring stage for non-client-facing jobs. The unit of observation is student-job-year, conditional on remaining in contention at each stage. All regressions include controls as well as job and year fixed effects. When firms conduct common screening rounds for multiple positions, each position is treated as a separate job. Bias-adjusted coefficients follow Oster (2019), setting $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$ (assuming unobservables match observables in selection intensity and allowing total explanatory power to exceed observed R^2 by 30%). For detailed discussion of these parameter choices, see Appendix E. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table OA.5: **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: This table shows caste differences in GPA between students included in and omitted from the earnings regression. "Omitted" comprises students who either deregistered from placement or registered but remain unplaced after the job fair, typically pursuing aftermarket opportunities, higher education, or entrepreneurship (Section 4.4). Column (1) includes only year fixed effects; Column (2) adds major and degree controls. Job fixed effects are excluded as they cannot be separately identified from the "omitted" coefficient for students without jobs. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

Table OA.6: **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 (OLS)	-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: This table reports within-job GPA differences across castes at each job search stage. Regressions control for major, degree as well as job and year fixed effects. Mean GPA is reported in raw GPA units for ease of interpretation (10 point scale). Bias-adjusted coefficients follow Oster (2019), setting $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$ (assuming unobservables match observables in selection intensity and allowing total explanatory power to exceed observed R^2 by 30%). For detailed discussion of these parameter choices, see Appendix E. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

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

	Dependent variable:					
	log (Avg. 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: This table reports earnings regressions for students with identifiable caste surnames at each job search stage. Each column represents a different stage: (1) application, (2) written test, and subsequent rounds. The dependent variable is log(average salary), where the average is taken over the posted salaries of all jobs a student remains eligible for at that stage (for the "Choice" column this collapses to the salary of the accepted job). All regressions include controls and year fixed effects. Job fixed effects are omitted as salaries are identical within jobs regardless of caste. Bias-adjusted coefficients follow Oster (2019), setting $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$ (assuming unobservables match observables in selection intensity and allowing total explanatory power to exceed observed R^2 by 30%). For detailed discussion of these parameter choices, see Appendix E. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

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

	Dependent variable:					
	log (Avg. 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						

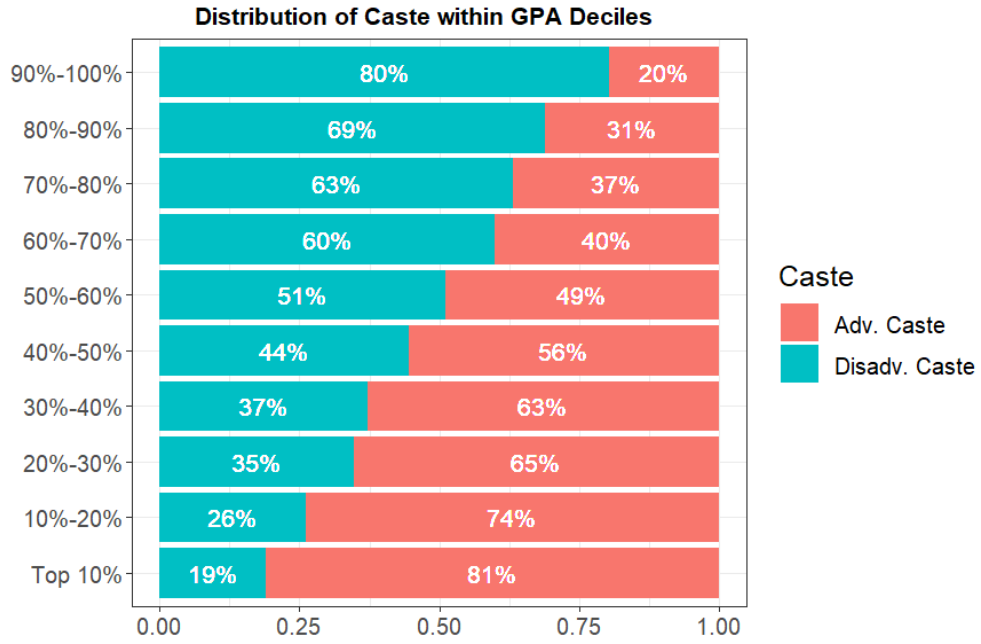
Notes: This table reports earnings regressions for students with non-identifiable caste surnames at each job search stage. Each column represents a different stage: (1) application, (2) written test, and subsequent rounds. The dependent variable is log(average salary), where the average is taken over the posted salaries of all jobs a student remains eligible for at that stage (for the "Choice" column this collapses to the salary of the accepted job). All regressions include controls and year fixed effects. Job fixed effects are omitted as salaries are identical within jobs regardless of caste. Bias-adjusted coefficients follow Oster (2019), setting $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$ (assuming unobservables match observables in selection intensity and allowing total explanatory power to exceed observed R^2 by 30%). For detailed discussion of these parameter choices, see Appendix E. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

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

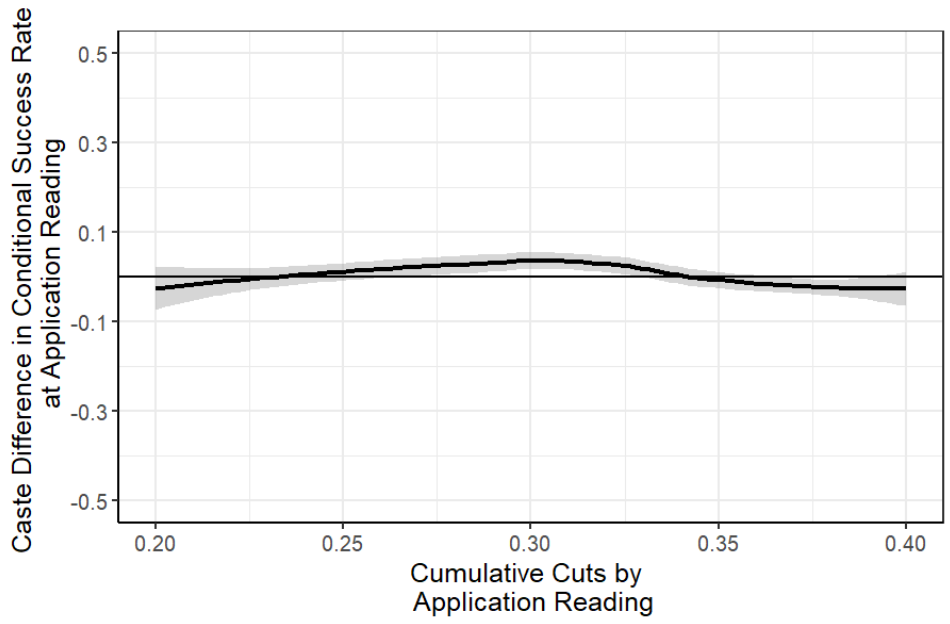
Notes: This table reports caste differences in GPA between job stayers and leavers. Column (1) includes year fixed effects only, Column (2) adds major and degree controls, and Column (3) additionally includes job fixed effects. Bias-adjusted coefficients follow Oster (2019), setting $\delta = 1$ and $R^* = 1.3 \times R$ (assuming unobservables match observables in selection intensity and allowing total explanatory power to exceed observed R^2 by 30%). For detailed discussion of these parameter choices, see Appendix E. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

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

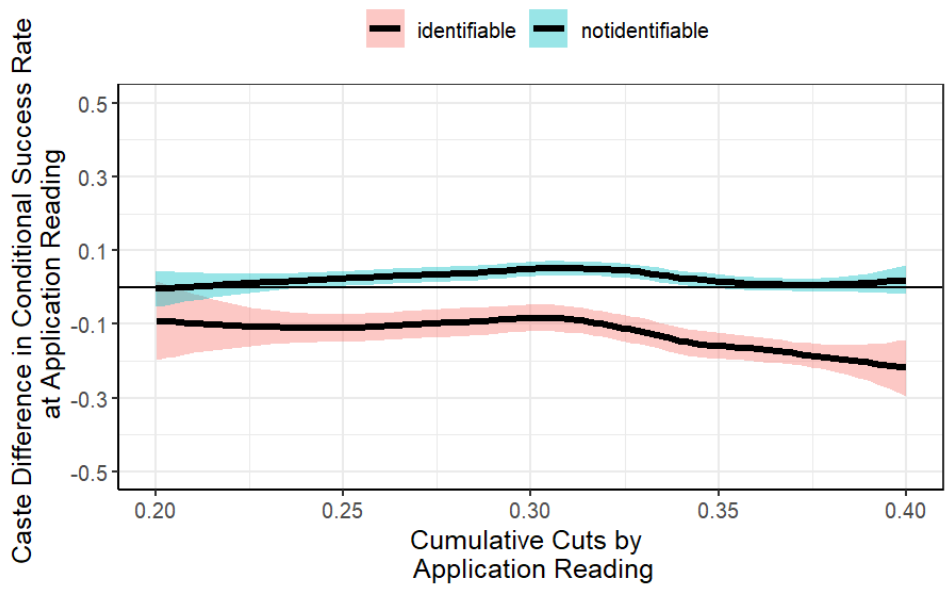


Notes: This figure shows the distribution of advantaged and disadvantaged caste students across GPA deciles. GPAs are normalized within degree-year.

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



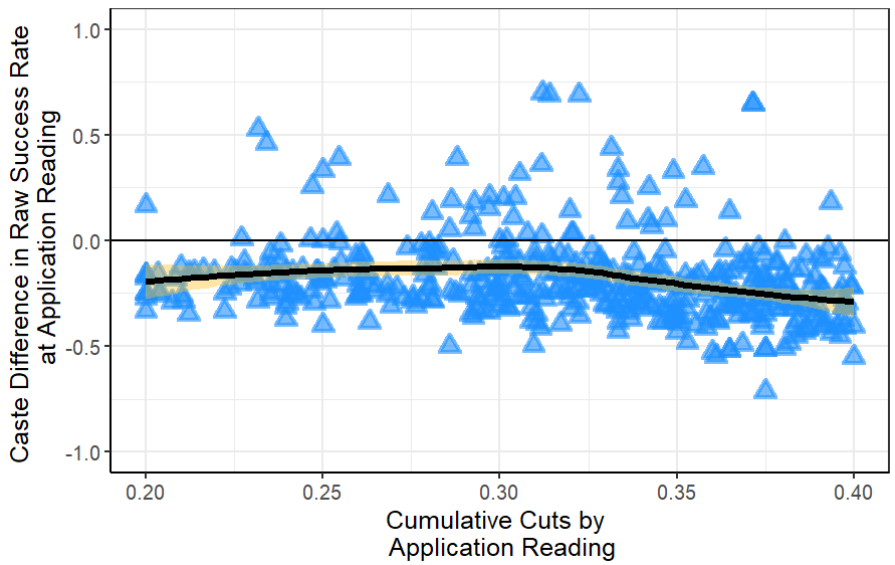
(a) Caste Penalty in Conditional Success Rates versus Cumulative Cuts by Application Reading in the Full Sample



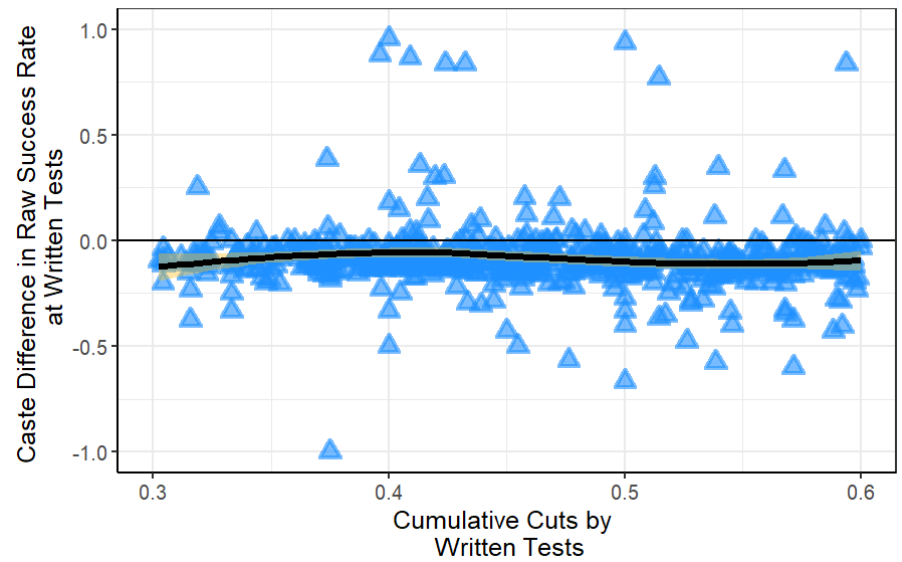
(b) Caste Penalty in Conditional Success Rates versus Cumulative Cuts by Application Reading in the Identifiable and Non-Identifiable Sample

Notes: This figure shows caste differences in success rates at the application reading round against cumulative selection cuts. Panel (a) uses the full sample; Panel (b) separates students with identifiable (red) and non-identifiable (blue) caste surnames. The x-axis shows the cumulative proportion of rejected applicants. Y-values represent the disadvantaged caste coefficient from job-specific linear probability models including controls (degree, major, GPA, entrance exam score, cohort year). Black lines show smoothed conditional means with 95% confidence intervals (shaded). The horizontal line at zero indicates no caste difference.

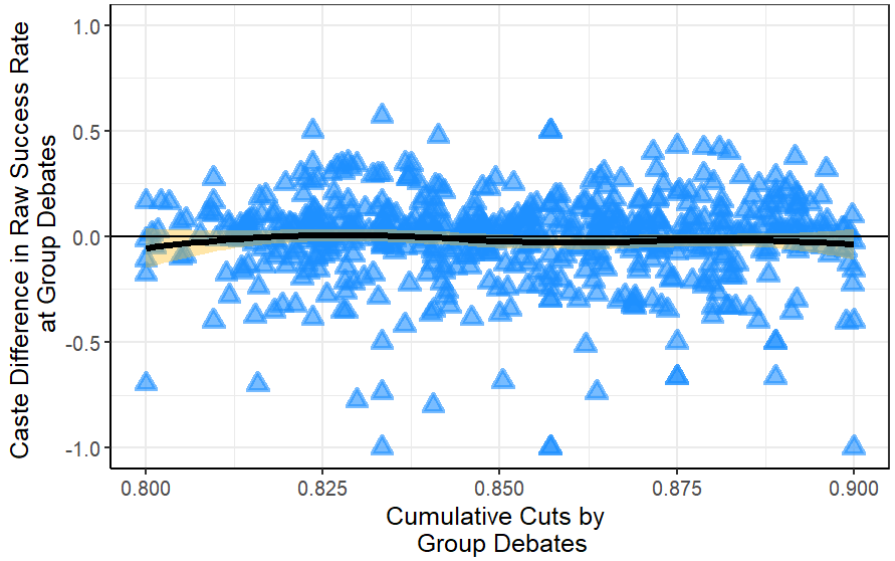
Figure OA.3: **Caste Penalty in the Uncontrolled Success Rate At Each Screening Round**



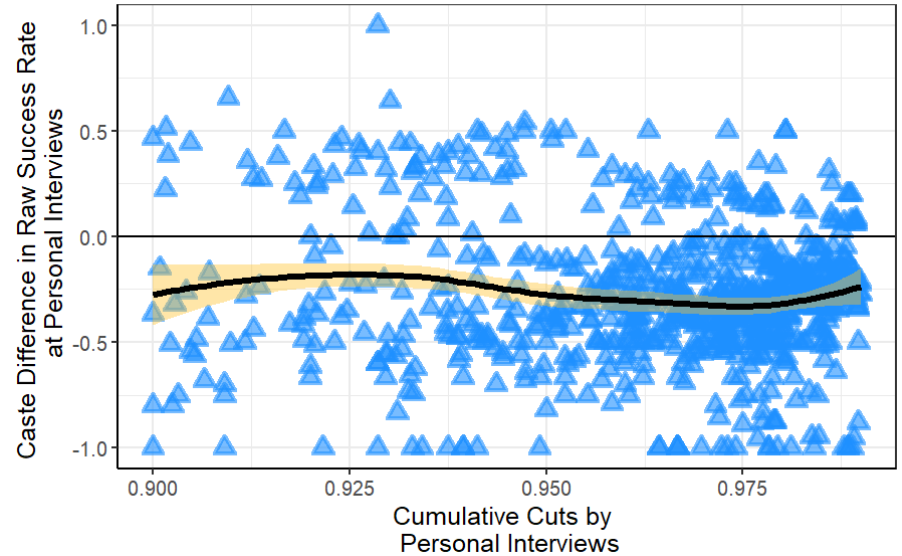
(a) Caste Penalty in Raw Success Rates versus Cumulative Cuts by Application Reading



(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: This figure shows raw (uncontrolled) caste differences in success rates against cumulative selection cuts at each screening stage. Each blue triangle represents a job. Black lines show smoothed conditional means with 95% confidence intervals (shaded). The x-axis shows the cumulative proportion of rejected applicants at each stage. The y-axis shows the raw caste gap in success rates. The horizontal line at zero indicates no caste difference.