Making the Elite: Class Discrimination at Multinationals

Soumitra Shukla*
Harvard Business School

Abstract

Class disparities persist in elite hiring, but where do they arise and why? Measuring class discrimination is difficult because social class is often inferred from overlapping, context-dependent markers—accent, schooling, networks, and cultural references—that do not map cleanly to a single indicator. India's codified caste categories provide a tractable proxy for cumulative social class, helping us identify where disparities arise in hiring and assess whether cultural-capital screening is a central channel of class bias in elite labor markets. Using screening data from over 1,000 jobs at multinationals recruiting at an elite Indian college, I track candidates across four rounds and find that subjective personal interviews that assess "cultural fit" account for over 95% of the caste hiring success-rate penalty. Two tests triangulate the mechanism. First, large language model analysis of nearly 1 million words of mock transcripts shows that personal interviews emphasize biographical markers of cultural capital—hometown, schooling, hobbies, family background, family structure, and parents' occupation—while prior subjective screening rounds such as group debates concentrate on abstract policy reasoning and discussion. Second, surname timing shows penalties at application screening for the 15% of students in my setting with caste-identifiable surnames, but only after "fit" interviews for the 85% with initially caste-ambiguous surnames. Taken together, these two tests indicate that for the surname-ambiguous majority, disparities emerge when interview content makes cultural-capital signals salient to evaluators and helps them discern caste, rather than from in-person cues alone in earlier face-to-face rounds. Disadvantaged-caste hires show lower attrition and higher promotion across job selectivity tiers and in both client-facing and non-client-facing roles, inconsistent with customer or statistical discrimination. Overall, the evidence points to cultural-capital screening as a central channel of class bias in elite multinational 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 Ramdas, 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)

"Every established order tends to produce ... the naturalization of its own arbitrariness." —P. Bourdieu, *Outline of a Theory of Practice* (1977)

1 Introduction

Despite his enormous wealth and polish, Jay Gatsby is dismissed in Fitzgerald's eponymous novel as "Mr. Nobody from Nowhere" (Fitzgerald, 1925). The line captures a common dynamic: biographical origins can continue to shape evaluation even when credentials appear strong. Elite hiring can work similarly when evaluators read cultural-capital markers—family background, linguistic polish, and social ease—as signals of merit. This paper quantifies how those inferences map into hiring penalties and when those penalties are consistent with bias.

Class disparities persist in elite employment outcomes (Laurison and Friedman, 2019; Reeves and Friedman, 2024; Stansbury, 2024), yet fundamental questions remain unanswered: How large are class-based hiring penalties, where in multi-stage hiring processes do they emerge, and why? Even within discrimination research, we rarely observe real-world hiring pipelines closely enough to pin down magnitude, location, and mechanisms (Altonji and Blank, 1999; Quillian et al., 2020)—and this knowledge gap is especially acute for class, with far less quantitative work than on race and gender (Stansbury, 2024).

This paper answers these questions using end-to-end recruitment data on campus hiring at an elite Indian college—an institution whose selectivity and global comparisons rank it alongside leading Ivy League institutions (Martellini et al., 2024). I trace candidates through four sequential screening rounds and find that final-round personal interviews[1] that assess "cultural fit" account for over 95% of the caste hiring success-rate penalty. Empirically, the estimated objects are caste-category penalties; I interpret them as informative about class-linked cultural capital in this setting. Converging evidence from surname timing, post-hire outcomes, and large language model (LLM) analysis indicates that these interviews screen biographical markers of cultural capital in ways difficult to reconcile with simple customer-discrimination predictions and with a broad class of accurate statistical-discrimination models, and consistent with a higher effective interview threshold for disadvantaged candidates.

Measuring class discrimination is difficult because class reflects intergenerationally transmitted advantage that often persists even as current income and credentials change (Clark,

[1] I use "personal interview" for the administrative stage in the hiring data and "HR-style fit interview" for the analogous transcript format in Section 6.4.

2014; Erikson and Goldthorpe, 1992; Reeves and Friedman, 2024). Unlike socioeconomic status (SES), which is a point-in-time snapshot of income and education, class captures *cumulative* economic, social, and cultural capital transmitted across generations (Bourdieu, 1984; Khan, 2011; Lareau, 2003; Weber, 1978a). Because this bundle is often only partially observed, researchers and evaluators typically infer class from combinations of cues rather than any single SES indicator. Many of these cues are interactional and embodied.[2] Pierre Bourdieu, one of the most influential theorists of class and stratification, argued that class reproduces itself not only through economic resources but through *cultural capital*[3]—markers such as accent, school pedigree, parental occupation, hobbies, and family background—that rarely map cleanly to a single variable and often surface through social interaction.

This creates a *double problem* for researchers: in most settings, class itself is measured imprecisely—should we use income, education, or parental occupation?—and even when a penalty is detected, the specific cue that triggered it remains unclear (Lamont, 1992; Rivera, 2012, 2015). India's caste system resolves both problems. Caste categories are administratively codified (Bayly, 2008; Dirks, 2001; Mandal Commission, 1980), capturing *cumulative* historical advantage and disadvantage spanning centuries, and are highly salient—98% of Indians identify with a caste group (Pew Research Center, 2021)—substantially reducing ambiguity about social position both for evaluators in interview rooms, who may use caste as a proxy for status, and for researchers trying to identify caste discrimination empirically. Ethnographic work on elite hiring suggests that subjective personal interviews probe biographical markers of class origin (Laurison and Friedman, 2019; Rivera, 2011), making personal interviews that assess "cultural fit" a plausible site where class discrimination operates.

Our setting provides unusual leverage for identifying and quantifying class discrimination. I study recruitment by multinational consulting, technology, and financial-services firms for their India operations. These jobs hire from an elite Indian public college, where hiring is centralized through the placement office and follows structured screening protocols. Because this institution operates under constitutionally mandated reservations, disadvantaged castes comprise roughly half of enrolled students (Madheswaran, 2008). Yet disadvantaged castes—who constitute nearly 70% of India's population—remain underrepresented in elite private-sector employment (Deshpande, 2011a; Jodhka, 2015). Since these first jobs shape career trajectories for decades in the world's most populous country (Subramanian, 2019), pinning down how much of this representation gap reflects discrimination—and *where and why* it arises—is a first-order question for long-run inequality.

[2] I use *embodied* to mean dispositions and styles acquired through socialization and expressed in speech, tastes, comportment, and interaction (Bourdieu, 1986; Goffman, 1959; Khan, 2011; Lareau, 2003).

[3] For a synthetic account of Bourdieu's framework, see Swartz (1997).

I observe over 1,000 jobs, each with four sequential rounds that move from relatively standardized to more subjective assessment: application reading, written aptitude tests, group debates, and personal interviews assessing candidate "fit." In the pooled sample, over 95% of the roughly 30-percentage-point caste penalty in hiring success rates emerges after personal interviews. Earlier rounds show negligible caste gaps after conditioning on observables (GPA, entrance exam scores, internships, major, and degree), while Section 6.4.2 documents an earlier penalty for the surname-identifiable subgroup. To address heterogeneity in unobserved candidate quality and hiring thresholds, I examine the full distribution of job selectivity (jobs that cut few versus many applicants before offers). The post-personal-interview penalty is similar across this distribution, including the most selective jobs, indicating that disparities are not confined to low-threshold positions or weaker candidate pools.

Our claim is not that fit interviews reveal no valuable information, but that they cannot justify the magnitude and pattern of disparities. Disadvantaged-caste hires have approximately 20% higher promotion rates and 28% lower attrition rates. A natural objection is inframarginality (Ayres, 2002): conditional on being hired, disadvantaged candidates may simply be unusually strong. If that were the full explanation, performance gaps should narrow in the most selective jobs where hires are closer to common margins. Instead, promotion and attrition advantages remain broadly stable across selectivity tiers, including the most selective jobs. This pattern is more consistent with disadvantaged candidates facing a higher effective interview threshold.

Three candidate explanations exist for the hiring penalty: accurate statistical discrimination, customer discrimination, and bias. Point estimates for post-hire outcomes move in the same direction in both client-facing and non-client-facing roles, making a standalone customer-discrimination account less plausible. The pattern is also hard to reconcile with accurate statistical discrimination: if employers held correct productivity beliefs (in means and/or variances), the post-hire advantage should tend to narrow as jobs become more selective.[4] Instead, disadvantaged workers outperform on retention and promotion across the selectivity distribution.

To triangulate this quantitative evidence, I use interviews I conducted with senior HR heads who oversee campus recruitment in India at Tata Consultancy Services, IRIS Software Group, Tata Motors, Aditya Birla Group, and Jindal South West Steel. These firms span IT services, enterprise software, manufacturing, steel, and diversified conglomerate businesses, with substantial operations in India and internationally. Across these firms, interviewees emphasize retention and promotion as key metrics signaling hire quality. I use this limited

[4]This attenuation prediction assumes monotone screening on expected productivity, reasonably well-behaved productivity distributions, and broadly stable task composition across selectivity tiers.

qualitative sample for institutional triangulation rather than stand-alone identification, but it strengthens confidence that the post-hire patterns are difficult to reconcile with accurate statistical discrimination. Taken together, this evidence is consistent with disadvantaged candidates facing a higher effective interview threshold, which can reflect taste-based discrimination, inaccurate beliefs about group productivity, or both; it does not by itself rule out every interview-specific latent trait.

An important mechanism question is what interviews are screening: directly observable caste labels, or class-like inference from biographical markers of cumulative social position. Put differently, do disparities for most candidates arise from overt labels, or when interviews make cultural-capital cues legible? This distinction maps to Bourdieu's account of class reproduction through cultural capital (Bourdieu, 1986): when explicit labels are weak, evaluators infer social position from biographical cues.

I use two complementary pieces of evidence to show that, in this setting, class-like screening through cultural-capital cues is an important channel through which evaluators *discern caste*, especially for candidates whose surnames are not initially caste-identifiable. First, because observing actual interview content at scale is rarely feasible (Quillian et al., 2020), I use LLM analysis of nearly 1 million words from public mock HR-interview and group-debate transcripts. These transcripts are used as a characterization of widely consumed preparation materials, not as direct transcripts of the focal employers' realized interviews. These YouTube preparation materials are widely used by candidates preparing for similar entry-level jobs in India and have cumulatively garnered hundreds of millions of views. The analysis shows what fit interviews plausibly probe in practice at two levels. At the *dominant-content level*, HR interviews are overwhelmingly about the candidate (89.7% coded "about speaker," versus 0.7% of group debates), whereas group debates are overwhelmingly about policy and society (99.3% coded "about society," versus 10.3% of HR interviews). At the *topic level*, HR interviews repeatedly elicit biographical markers of cultural capital—hometown, schooling, hobbies, family background, and parental occupation—while group debates instead center on policy topics such as economic policy, education policy, governance and democracy, technology, and artificial intelligence.

The same interviews I conducted with senior HR heads provide a triangulation check for whether the transcript evidence reflects actual screening content: interviewees describe "fit" rounds as conversational and open-ended, with biographical questions about family background and schooling often arising in practice. As one interviewee put it, "In the HR round, we do ask for family backgrounds and other things." Another emphasized how social background enters communication assessments, noting that some candidates

come from families "where their parents were working as domestic help or driver," so evaluators "can't expect them to read encyclopedias and Enid Blyton."[5] Taken together, these accounts suggest that communication and background probes are folded into broader cultural-fit assessments. While such assessments may in principle aim to improve matching, triangulating these interviews with linked post-hire outcomes shows disadvantaged-caste hires have lower attrition and higher promotion—the same downstream metrics interviewees emphasize as indicators of hire quality. This alignment between interview accounts, transcript evidence, and downstream outcomes strengthens the credibility of using mock transcripts to characterize interview content in this setting and is less consistent with a pure productivity-screening interpretation.

The second piece of evidence ties interview content to caste revelation. While some surnames are caste-identifying, most are ambiguous in urban-educated settings because migration, "Sanskritization,"[6] and regional naming conventions blur name-based inference (e.g., Jodhka, 2015; Mamidi, 2011; Srinivas, 1966; Wiltshire, 2020). In this setting, only about 15% of students have caste-identifiable surnames. Among this 15%, penalties emerge at application reading; for the 85% with caste-ambiguous surnames, disparities emerge only after personal interviews. This timing pattern indicates that penalties plausibly arise when interview interactions make caste legible. Combined with the LLM evidence, this points to cultural-capital screening as a central channel through which evaluators discern caste in "fit" interviews.

Why might such bias persist? My interviews with senior HR heads point to two organizational frictions. First, campus hires are recruited into central pools rather than matched to specific managers, creating weak interviewer-supervisor alignment ("the placements are so random, that it's almost like lottery"; "there is no line manager to student alignment"). Second, recruiters are evaluated mainly on process metrics (targets, turnaround time, join rates), not downstream hire performance. Interviewees also noted that post-hire outcomes are difficult to trace back to recruitment because team assignment, onboarding, and workplace dynamics intervene. This separation and weak feedback may reduce incentives to correct biased fit screening.

Scholars disagree on whether caste is best understood as analogous to race, ethnicity,

[5] Enid Blyton was a widely read twentieth-century British children's author; in this context, the reference functions as shorthand for early English-language cultural exposure that interviewers may read as part of "fit."

[6] "Sanskritization" refers to mobility processes in which lower-status *jāti* (localized endogamous caste) groups adopt names, customs, and practices associated with higher-status groups to claim upward status mobility (Srinivas, 1966).

class, or as a distinct institution that partially overlaps all three.[7] In this paper, the class framework is used in a *measurement-focused, not ontological* sense. Three observations support this framing. First, caste penalties are most pronounced at interview stages that elicit biographical markers—especially for the roughly 85% of candidates with caste-ambiguous surnames—rather than at first visual exposure in group debates. Second, the specific markers interviews extract—hometown, schooling, parental occupation—are those that leading sociological theories of stratification treat as constitutive of class (Blau and Duncan, 1967; Bourdieu, 1986; Erikson and Goldthorpe, 1992; Lareau, 2003; Weber, 1978a). Third, in the Supreme Court of India's landmark decision in *Indra Sawhney v. Union of India* (1992), the Court states that "[a] caste can be and quite often is a social class in India" (Supreme Court of India, 1992). The *Mandal Commission*—constituted to identify socially and educationally backward groups for affirmative-action policy—institutionalized the "Other Backward *Classes*" (OBC) category and defined "disadvantage" through bundled indicators (occupation, education, housing, marriage patterns, and social standing) that capture cumulative social position (Mandal Commission, 1980). Operationally, the Commission used an eleven-indicator weighted scoring rule (social indicators weighted 3, educational 2, economic 1; maximum score 22) and classified groups scoring at least 11 as backward (Galanter, 1984; Mandal Commission, 1980). These three observations justify treating caste as a high-signal measure of cumulative class position, clarifying what the empirical results identify and grounding the contribution claims that follow.

Related literature. This paper contributes to discrimination research by providing, to my knowledge, the first large-scale evidence from major multinationals that identifies and quantifies cultural-capital screening as a mechanism of bias within a real hiring pipeline. This is new because discrimination research generally does not observe multi-stage hiring deeply enough to recover mechanism salience or adjudicate whether this channel reflects efficient screening or bias (Bertrand and Duflo, 2017; Small and Pager, 2020). Using stage-by-stage outcomes alongside surname timing, mock-transcript evidence on interview content, and post-hire comparisons, I show interview-stage penalties are difficult to reconcile with simple customer-discrimination predictions and with broad accurate statistical-discrimination accounts, and are consistent with a higher effective interview threshold for disadvantaged

[7] This debate is longstanding. Ghurye observed that "we do not possess a real general definition of caste," and that any single definition is difficult because of the phenomenon's complexity (Ghurye, 1969). Inden similarly notes the absence of a universally accepted definition (Inden, 2001). Competing traditions emphasize ritual hierarchy (Dumont, 1966), colonial-administrative reification (Bayly, 2008; Dirks, 2001), ethnicity-like *jāti* boundaries (Béteille, 1996), and socioeconomic stratification in labor and status terms (Bourdieu, 1986; Deshpande, 2011b; Weber, 1978a).

candidates. Over 95% of penalties emerge at personal interviews—a margin canonical audit and correspondence designs cannot observe (Quillian et al., 2020). This implies that upstream anonymization policies (for example, CV blinding) may reduce initial gaps yet still miss substantial downstream discrimination generated in interaction-heavy interviews.

Second, the paper contributes to class-discrimination research, where evidence remains thinner than in race and gender (Stansbury, 2024). A central challenge is that cultural-capital screening is difficult to identify and quantify: the relevant cues are interactional, multidimensional, and typically unobserved in administrative hiring data (Goffman, 1959; Savage, 2015). Researchers therefore rely on point-in-time proxies—income, education, and parental occupation—that imperfectly capture cumulative intergenerational position (Blau and Duncan, 1967; Erikson and Goldthorpe, 1992). Canonical sociological scholarship conceptualizes cultural capital as embodied, interactional, and relational, and links it to class reproduction and evaluative boundary-drawing (Bourdieu, 1984, 1986; Lamont and Lareau, 1988; Lareau, 2003; Savage et al., 2015; Swidler, 1986). But that tradition is typically not designed to recover mechanism salience in multi-stage labor-market screening: it rarely quantifies how much of observed hiring inequality this channel generates, or where in the pipeline it binds. Building on this literature, I add a quantitative layer by combining caste as a measurement device for cumulative social position (Bayly, 2008; Mandal Commission, 1980), stage-wise decomposition of where penalties emerge, and LLM analysis of interview content. This integrated design directly addresses that gap by identifying when biographical markers become legible, quantifying how strongly penalties concentrate within the pipeline, and connecting those patterns to downstream outcomes.

Third, the paper contributes to personnel and organizational economics by providing field evidence on the allocative consequences of discretionary fit interviews in high-stakes hiring. Research on person-organization fit and enculturation shows that organizations value cultural matching (Chatman, 1991; Goldberg et al., 2016; O'Reilly et al., 1991), while personnel research continues to debate the predictive value of unstructured interviews relative to structured or mechanical approaches (Benson and Shaw, 2024; Highhouse, 2008; Hoffman and Stanton, 2024; Kuncel et al., 2013; Sackett et al., 2022). This creates a core tension: enculturation can facilitate coordination, but discretionary fit screening can also reward social similarity rather than job-relevant capability (Srivastava et al., 2018). The evidence here speaks directly to that tension. In this setting, the stage with the greatest discretionary fit evaluation is also where disparities concentrate, even though post-hire outcomes do not indicate weaker performance among disadvantaged hires. This pattern is consistent with fit-intensive screening placing weight on social-similarity criteria that are imperfectly aligned with productivity-relevant

allocation, consistent with broader evidence that subjective evaluation systems can distort allocation and advancement (Hoffman et al., 2018; Sarsons, 2017).

Fourth, the paper contributes to mobility and stratification debates by showing how downstream hiring discretion can offset upstream equalization. Admissions-side policies can broaden access to elite colleges (Karabel, 2005; Stevens, 2009), including affirmative action in this setting (Bagde et al., 2016; Madheswaran, 2008; Yenmez and Sönmez, 2022), but evidence from elite labor markets shows that the conversion of elite education into top jobs and incomes remains unequal by social background (Jodhka and Naudet, 2019; Zimmerman, 2019; Michelman et al., 2022). In these data, this disconnect is concrete: representation is broadened at college entry through government-mandated affirmative-action policies that increase disadvantaged-caste representation among graduates (Madheswaran, 2008), yet disparities remain concentrated at the final interview stage of the hiring pipeline. The pattern is consistent with frameworks in which cultural transmission and group-based social interactions sustain inequality despite formal access expansion (Bisin and Verdier, 2025; Cassan et al., 2021; Durlauf, 2006). The implication is that admissions-side equalization alone may not translate into durable mobility without changes in downstream hiring.

Fifth, the paper contributes to caste-discrimination research, where labor-market evidence remains relatively limited compared with adjacent work on educational affirmative action and broader caste stratification (Bagde et al., 2016; Bayly, 2008; Bertrand et al., 2010; Jodhka, 2015). Notable labor-market exceptions include Asher et al. (2024), Cassan et al. (2021), Deshpande (2011a), and Munshi (2019). Relative to this literature, the contribution here is granular pipeline evidence that quantifies where penalties arise within hiring and identifies cultural-capital screening in subjective interviews as a quantitatively salient mechanism.

2 Background: India's Administrative Caste Categories as a Measure of Social Class

Scholars debate whether caste is best conceptualized as race, ethnicity, class, or a distinct institution that overlaps all three (Bayly, 2008; Dirks, 2001; Dumont, 1966; Ghurye, 1969). I use an explicitly measurement-focused, not ontological, framework: in this hiring setting, India's administratively codified caste categories provide a workable proxy for cumulative social position. The goal is to justify why a class framework is analytically appropriate for the mechanism studied here—interview-stage biographical screening—while leaving fuller conceptual detail to Appendix A (Table A.1 provides a comprehensive taxonomy of caste,

class, and race across mechanism-relevant dimensions for discrimination research).

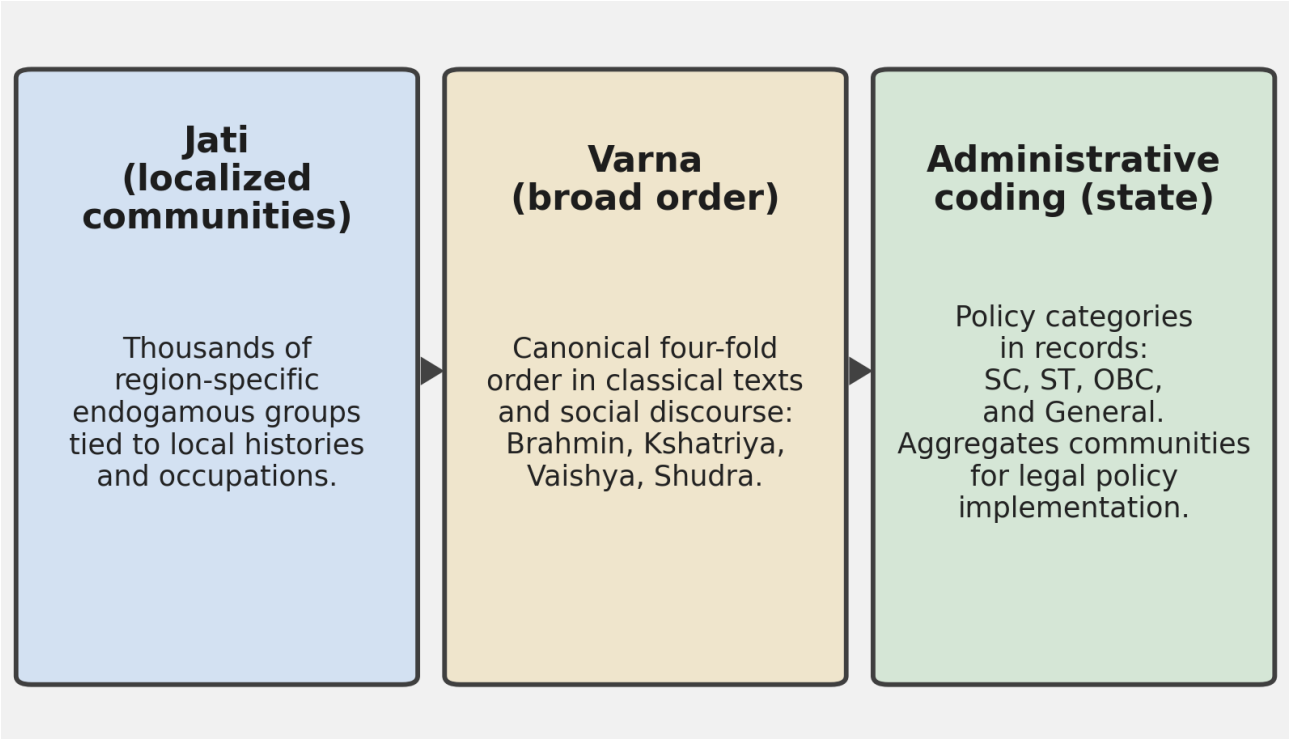
2.1 India's Administrative and Legal Framework

The English term "caste" conflates distinctions that Indian scholarship keeps separate: *jāti* refers to localized endogamous groups, while *varna* denotes broader classificatory order (Basham, 1954; Béteille, 1996; Cohn, 1987). In classical usage, the *varna* schema names four broad orders—*Brahmin*, *Kshatriya*, *Vaishya*, and *Shudra*—whereas policy categories in modern India are constructed through administrative aggregation rather than one-to-one mapping onto this four-fold scheme (Bayly, 2008; Dirks, 2001). Many communities now coded as SC are historically associated with Dalit/*avarna* groups treated as outside the four-*varna* order, and many communities coded as ST are tribal groups likewise outside that schema; OBC coding is heterogeneous and not reducible to a single *varna* bin (Bayly, 2008; Cohn, 1987; Dirks, 2001; Galanter, 1984).

> "The term 'caste' was applied indiscriminately to both varna or class, and jati or caste proper. This is a false terminology." (Basham, 1954)

Historically, the key administrative bridge is explicit: late-colonial "Depressed Classes" terminology was formalized into the statutory category "Scheduled Castes" through the Government of India (Scheduled Castes) Order, 1936 under the Government of India Act, 1935 (His Majesty in Council, 1936; Bayly, 2008; Galanter, 1984). In late-colonial usage, "Depressed Classes" referred to communities treated as socially excluded and historically subject to untouchability and related civic disadvantage. Early representation politics in Mysore included petitions to Maharaja Krishnaraja Wadiyar IV and fed into the Miller Committee process on community representation in public service (Miller, 1919; Rajya Sabha Committee on Personnel, Public Grievances, Law and Justice, 2005). The committee, chaired by Leslie C. Miller, was explicitly appointed to examine steps for "adequate representation of communities in the public service," making it an early institutional template for group-based representation policy (Miller, 1919). The term "Scheduled Castes" comes directly from these official schedules. Figure 1 summarizes the measurement sequence used in this paper: localized *jati* identities map imperfectly into broader *varna* discourse, while empirical coding uses the modern administrative categories (SC/ST/OBC/General) that aggregate communities for policy implementation. Appendix A, Figures A.1–A.3, show archival schedule evidence for this terminology.

Figure 1: **From *Jati* and *Varna* to Administrative Categories**



Notes: This conceptual schematic is a measurement aid for the empirical design. It distinguishes localized *jati* communities from broader *varna* discourse and then maps to the administratively observed categories used in the data (SC/ST/OBC/-General). The schematic highlights that SC and ST coding includes communities commonly treated as outside canonical *varna* classification, and does not claim a one-to-one historical mapping from *varna* to modern policy categories. Historical statutory schedule evidence is shown in Figures A.1–A.3.

For measurement, the key point is administrative aggregation: Scheduled Castes (SC, communities historically subjected to untouchability), Scheduled Tribes (ST, indigenous communities historically marginalized by geography and state reach), and Other Backward Classes (OBC, groups classified as socially and educationally backward under post-independence policy and Mandal Commission criteria) extend beyond any literal varna mapping but are operationalized through a common framework of social and educational disadvantage. In practice, states classify and aggregate many localized communities into these categories for policy implementation, including groups not cleanly representable within a ritual four-varna frame (Bayly, 2008; Cohn, 1987; Dirks, 2001). This administrative approach is codified in India's constitutional and policy architecture, which explicitly uses class language ("backward classes")[8]. Empirically, this study uses this administrative layer—not inferred *varna* status and not community-by-community *jati* coding—because that is the legally verified classification observed in college admission records and then linked to hiring outcomes.

[8]For constitutional provisions (including Articles 340, 341, and 342) and statutory implementation details, see Government of India (1950); Galanter (1984); Mandal Commission (1980).

The Mandal Commission is central because it translates the constitutional language of "backward classes" into an operational classification rule. It identified OBC groups using an eleven-indicator composite (four social, three educational, and four economic indicators), weighted social and educational criteria more heavily than economic criteria, and classified groups scoring at least 11 out of 22 as backward (Galanter, 1984; Mandal Commission, 1980). This multidimensional, weighted structure treats disadvantage as a bundled social position rather than income alone.

The landmark judgment in *Indra Sawhney v. Union of India* (1992)—often referred to as the "Mandal case" because it adjudicated implementation of the Mandal framework—provides the key judicial articulation. The Supreme Court stated:

> "[A] caste can be and quite often is a social class in India." (Supreme Court of India, 1992)

The ruling pairs this with social-and-educational backwardness requirements and creamy-layer exclusion.[9] For this paper, the implication is identification: Mandal's indicator structure and *Indra Sawhney*'s doctrine jointly justify treating SC/ST/OBC categories as codified proxies for *cumulative* disadvantage, and the hiring pipeline then reveals when disparities emerge.

2.2 Academic Scholarship on Caste and Class

Academic scholarship does not imply a single definition of caste. Ghurye's classic caution is still apt:

> "We do not possess a real general definition of caste. It appears to me that any attempt at definition is bound to fail because of the complexity of the phenomenon." (Ghurye, 1969)

One strand emphasizes ritual hierarchy and endogamy (Dumont, 1966; Srinivas, 1966); a second emphasizes administrative legibility and state codification (Bayly, 2008; Cohn, 1987; Dirks, 2001); a third treats caste as materially grounded stratification tied to occupation, education, marriage markets, and mobility (Béteille, 1965, 1996; Cassan et al., 2021; Deshpande, 2011a; Jodhka, 2015). This paper uses the third strand in a measurement-focused way while not denying the others. As Béteille puts it:

> "At the level of varna, the old classical scheme of four orders ... has always been more of a class than an ethnic division." (Béteille, 1996)

[9] "Creamy layer" excludes relatively better-off OBC households from reservation benefits in public-sector admissions and jobs while group classification remains OBC (Galanter, 1984; Supreme Court of India, 1992). In this private-sector setting, OBC creamy-layer candidates still face substantial interview penalties (see Section 2.3).

Weber's canonical status-group account similarly links social honor to social and economic closure, and treats caste as the limiting case of that closure:

> "Where the consequences [of status distinctions] have been realized to their full extent, the status group evolves into a closed 'caste'." (Weber, 1978a)

Leading class and stratification scholarship emphasizes that culturally encoded dispositions can transmit advantage in evaluative settings. In Bourdieu's framework, cultural capital has multiple forms; the channel studied here is primarily *embodied*. In this paper's usage, *embodied cultural capital* means durable interactional dispositions—for example accent, communicative style, conversational ease, and familiarity with elite cultural references—formed through family and schooling and made legible in interviewer-candidate interaction (Bourdieu, 1986; Khan, 2011; Lareau, 2003; Lamont, 1992). Related work on caste and stratification links caste-linked social position to networks, occupations, education, and mobility (Béteille, 1965; Cassan et al., 2021; Deshpande, 2011a; Jodhka, 2015; Munshi, 2019). In this setting, those scholarship-based expectations map directly to the evidence: personal interviews probe biographical and interactional cues, while group debates place much greater weight on abstract policy reasoning; together with surname timing, this pattern aligns with inference through embodied cultural capital and is harder to reconcile with purely phenotypic identification (Bourdieu, 1986; Kraus et al., 2017; Lamont, 1992; Omi and Winant, 2014).

This framing is also consistent with persistence evidence: class systems are formally more mutable than caste categories, yet social rank remains highly persistent across generations (Clark, 2014; Durlauf, 2006). Appendix A provides the fuller conceptual discussion.

2.3 Operational Coding of Caste in Affirmative Action Framework

Operationally, I code SC/ST/OBC candidates as *Disadvantaged* and General/Unreserved candidates as *Advantaged*, matching administrative categories while interpreting them as proxies for cumulative social position in private-sector hiring. Affirmative action in India applies to public education and government jobs, not private-sector multinational hiring (Galanter, 1984); the mechanism documented here therefore operates where legal reservation protections do not directly constrain employer screening. Disaggregated estimates support this coding: even OBC candidates above creamy-layer thresholds face substantial interview penalties, inconsistent with income-alone explanations. Appendix B reports category-disaggregated stage outcomes, job-quality patterns, and the OBC creamy-layer comparison.

Appendix C reports India-facing DEI evidence from firms recruiting at elite Indian campuses: 81% of sampled firms have no substantive DEI page, and among firms with substantive

DEI pages, caste appears in only 5.1% of India-headquartered pages and 1.0% of foreign-headquartered pages. This low formal-policy visibility does not prove the absence of internal informal norms. But India-specific qualitative scholarship reports substantial employer resistance to caste-conscious affirmative action and continued reliance on merit/fit language in corporate hiring discourse (Jodhka, 2008; Jodhka and Newman, 2007; Thorat and Newman, 2010).

3 Key Terms and Institutional Setting

This section defines key terms and describes the main institutional setting: an on-campus job fair at an elite Indian college (a single institution, anonymized per data agreement). This job fair is representative of elite hiring in India—nearly all top-ranked colleges use similar standardized placement procedures, graduates predominantly pursue elite private-sector positions, elite employers recruit almost exclusively from the same set of campuses, and student participation rates are near-universal. Appendix D documents this representativeness evidence.

3.1 Definitions

Five key terms are used throughout the analysis:

a. **Disadvantaged castes.** Scheduled Castes (SCs), Scheduled Tribes (STs), and Other Backward Classes (OBCs), as defined in Section 2.3.

b. **Advantaged castes.** Castes not classified as SC, ST, or OBC—commonly referred to as "General Category" in India's administrative framework.

c. **Job.** A firm-designation pair (e.g., Google software engineer).

d. **Elite jobs.** Entry-level positions at multinationals, major national corporations, or unicorn startups (privately-held startups valued at over $1 billion).

e. **Caste-based affirmative action.** Constitutional reservations requiring government jobs and educational institutions to allocate up to 50% of seats or positions to SC, ST, and OBC candidates combined at central institutions (some states exceed this ceiling); these reservations do not apply to private sector positions (see Section 2.3).

3.2 The Job Fair

The elite college's job fair occurs *entirely on campus* and comprises five phases:

a. **Pre-Placement.** Firms deliver pre-placement talks between June and mid-August, gauging student interest and finalizing job roles.

b. **Return offers from internships.** By late August, firms extend pre-placement offers to summer interns. *Unlike in the U.S., these play a minor role*: fewer than 5% of students accept and exit the placement process.

c. **Student and employer registration.** Students register by late August; employers submit job roles, compensation, and vacancies by early September. Compensation is *non-negotiable*: the placement office verifies offer letters match advertised pay, and non-compliance can result in restrictions on future campus visits. Such rules and timelines are common across elite Indian colleges and are publicly available on their websites (e.g., IIT Delhi Office of Career Services; Delhi University Placement Cell).

d. **Placement.** Starting mid-September, students apply for jobs. Employers screen applications, administer written aptitude tests and group debates (September–December), then conduct final-round personal interviews (December–January).

e. **Post-Placement ("aftermarket").** From late January until graduation, the placement office runs an "aftermarket" for students without offers. Aftermarket employers—typically lower-paying startups—use less standardized recruitment.

4 Data and Sample

The analysis draws on administrative records from an elite Indian college's centralized job portal spanning four placement cycles (exact years anonymized). The dataset includes 4,164 students pursuing technical degrees competing for over 1,000 positions in software engineering, quantitative consulting, data analytics, and manufacturing at multinational firms.

Data Source. Through the college's centralized placement portal, students upload standardized profiles (GPA, entrance exam scores, coursework, internship experience) and employers record and upload shortlists after each screening stage, ensuring complete documentation from application through final offer. Registered students are mandated to participate in all

four screening rounds, and offer-letter salaries are verified by the placement office against the non-negotiable compensation schedules filed by firms.

Student Sample. The 4,164 students pursue four technical degrees: Bachelor of Technology (B.Tech.), Dual Degree, Master of Technology (M.Tech.), and Master of Science (M.S.). Admissions are *exclusively* based on standardized national-level entrance examinations with caste- and major-specific cutoffs—no subjective components enter the process. Approximately 90% of students are male, typical of elite technical colleges in India (Datta, 2017). Appendix O, Table O.1 provides the full breakdown by degree.

Caste Measurement and Representation. *Caste-based affirmative action* in college admissions creates approximate parity in the overall sample (48% disadvantaged, 52% advantaged). Appendix O, Table O.1 shows near-parity in most degree programs, though M.S. students show lower disadvantaged-caste representation (27%) due to smaller cohort sizes and distinct admission pathways (with small n, realized shares can deviate from the approximately 50% quota). India's reservation system mandates that a fixed proportion of seats be allocated to historically disadvantaged castes (Scheduled Castes, Scheduled Tribes, and Other Backward Classes), with caste status verified through government-issued certificates that are mandatory for claiming reservation benefits and cross-checked against administrative records (Galanter, 1984). All students declare caste at admission; statutory certificates are required only for reserved-seat claimants. For OBC students, caste is verified via government certificates for all claimants, while the creamy-layer/non-creamy-layer distinction is determined by income documentation at admission. I refer to OBC creamy-layer students as OBC-CL.

We treat OBC-CL students as disadvantaged despite their ineligibility for reservation benefits. Appendix B shows that OBC-CL candidates face the largest interview penalties, consistent with a Bourdieusian framework where cumulative cultural and social capital—not current income alone—determines social position.

Researcher Observes State-Verified Caste Categories; Recruiters Do Not. These state-verified caste records are available to me in the college placement data, but recruiters are *not directly provided* candidates' caste categories during screening. This distinction is central to interpreting the surname timing evidence (Section 6): for candidates with caste-ambiguous surnames, recruiters must infer caste from biographical markers rather than observing it directly. The 15% surname-identifiable share should therefore be interpreted as a lower bound on pre-interview caste legibility, because recruiters may also infer caste from other cues (for

example first name, hometown, schooling language, accent, or communication style).

Summary Statistics for Students. Disadvantaged castes score substantially lower on GPA and entrance exams,though both groups span all performance deciles (Appendix N, Figure N.1). All analyses control for these differences. Despite academic gaps, caste shows no statistically significant association with internship sector: among undergraduates, distributions are comparable across castes (25% technology, 37% consulting, 38% manufacturing); similar results hold for postgraduate degrees (not shown).[10] Appendix O, Table O.2 provides detailed breakdowns by caste and degree.

Summary Statistics for Jobs. The dataset includes over 1,000 private-sector jobs, mostly offered by multinationals headquartered in the U.S. or Europe. Around 75% of roles are non-client-facing (e.g., software engineering) and 25% client-facing (e.g., quantitative consulting). Public-sector roles (about 2%) are excluded due to distinct salary structures. Appendix O, Table O.3 provides detailed job characteristics.

Career Outcomes Data. Tenure and promotions are obtained for over 98% of placed students through alumni records and employment platforms. Appendix E reports coverage benchmarks.[11] Promotions are defined as movements to higher job titles within the same firm, validated by salary increases and industry benchmarks (e.g., Wilmers et al., 2025).

Data Limitations. Two limitations merit note: the absence of detailed résumés beyond portal data and the lack of screening-round performance scores. The former is mitigated by comprehensive portal data (GPA, entrance scores, internships); the latter is addressed using variation in job selectivity (Section 6.2).

5 Institutional Features of the Hiring Process

This section describes three institutional features of the hiring process: the low-cost application system, the standardized four-stage recruitment pipeline, and key features (standardized compensation, internship policies, and recruiter composition) that shape interpretation of the empirical patterns.

[10] Such patterns likely reflect that internship return offers are rare (Section 3.2), making internship choices primarily driven by student sector preferences rather than caste-related pay or opportunities.

[11] The missing 2% show no systematic caste differences in observable characteristics.

Applications. Because applying to jobs involves only clicking on the company name in the centralized portal, students apply broadly within eligibility constraints.

Four-Stage Recruitment Pipeline. Once applications are submitted, recruitment progresses through four sequential stages *for all jobs*: (1) *Application screening*, where employers typically use GPA cutoffs to filter résumés; (2) *Written aptitude tests*, standardized assessments of technical and analytical skills; (3) *Group debates*, structured discussions involving 20–30 students divided into two teams with uniform prompts (employers run 4–5 concurrent sessions to accommodate all candidates); and (4) *Personal interviews*, around 30-minute final-round evaluations emphasizing interpersonal skills and organizational and cultural fit.

In this market, the final round is typically an HR-style fit interview distinct from earlier technical screens. Recruiters observe standardized profile fields (name, GPA, entrance exam rank, school marks, internships) from the application stage onward; group debates add real-time communication cues; and personal interviews then add direct biographical probing (for example hometown, schooling trajectory, family background, and parents' occupation), which is where caste salience can intensify for surname-ambiguous candidates.

Substantial selectivity operates at each stage: on average, jobs receive roughly 209 applications, with approximately 125 advancing to written tests, 102 to group debates, 26 to interviews, and about 6 receiving offers. Appendix O, Table O.4 reports these stage-by-stage counts.

Standardized, Non-Negotiable Compensation. Salaries and benefits are *standardized, non-negotiable, and uniform* for each job role, independent of caste, gender, or major—ensuring any earnings gaps reflect differential job access rather than wage-setting idiosyncrasies.

Internship Return Offer Policy. Students accepting internship return offers are withdrawn from the placement cycle by the placement office and excluded from analysis.

India-Based Positions and Indian Recruiters. About 99% of positions are India-based; consequently, recruiters and interviewers are almost always Indian hiring managers.

6 Caste Penalties Emerge at "Fit" Interviews

Disadvantaged caste candidates face a substantial penalty in hiring outcomes, and over 95% of the caste penalty in stage-specific success rates emerges at personal interviews—the "fit"

assessment. First, Section 6.1 shows that caste gaps open sharply at interviews, not earlier stages in the pooled sample. Second, Section 6.2 shows that this interview penalty is flat across the full distribution of job selectivity—from jobs that cut relatively few applicants to jobs that cut many—inconsistent with cumulative ability deficits being revealed only at interviews.

Third, Section 6.3 shows that disadvantaged hires exhibit higher promotion rates and lower attrition across selectivity tiers and in both client-facing and non-client-facing roles; interviews I conducted with senior HR heads (a limited qualitative sample) inform this interpretation by indicating that retention and promotion are central measures of hire quality. Section 6.3 also uses a stylized model to formalize why stable post-hire advantages across selectivity tiers are difficult to reconcile with cumulative productivity deficits. Fourth, Section 6.4 combines LLM transcript evidence with surname timing to show that penalties track caste revelation through biographical probing of cultural-capital markers; the same interviews I conducted indicate that personal interview rounds often probe family background and schooling. Fifth, Section 6.5 adjudicates between accurate statistical discrimination, customer discrimination, and bias, and uses interviews I conducted with senior HR heads to provide suggestive qualitative evidence on organizational frictions that may allow biased screening to persist. Across these sections, qualitative triangulation draws on interviews I conducted with senior HR heads at major multinationals; full theme-ordered excerpts are reported in Appendix F.

Identification interpretation. The selectivity and post-hire tests rely on three maintained assumptions (formalized in Appendix I and Appendix H): (i) across jobs, higher cumulative cuts proxy tighter screening thresholds rather than caste-specific resorting into fundamentally different applicant-quality distributions; (ii) promotion and first-job leaving are informative post-hire outcomes for match quality, conditional on controls and fixed effects; and (iii) the evidence rules out broad pre-interview cumulative-deficit stories that should surface in high-cutoff jobs, but does not by itself rule out every interview-specific trait first observed in round 4.

6.1 Caste Penalties in Hiring

Success Rates. The sample for success rate analyses comprises 4,005 students—96% of the 4,164 graduating cohort—who participated in the centralized placement process. Binary advancement outcomes (advance/not advance at each stage) are defined for *all participants*, regardless of whether they ultimately secure employment.

To estimate caste differences in success rates *at each stage*, I estimate the following linear probability model:

$$\text{Advance}_{ijt} = \alpha + \beta \cdot \text{Disadvantaged Caste}_i + \mathbf{X}_i'\boldsymbol{\gamma} + \delta_j + \theta_t + \epsilon_{ijt}, \tag{1}$$

where Advance_{ijt} is an indicator equal to 1 if student i advances to the next stage for job j in cohort-year t; $\text{Disadvantaged Caste}_i$ is an indicator for belonging to a disadvantaged caste (SC/ST/OBC); δ_j are job fixed effects; θ_t are cohort-year fixed effects; and $\mathbf{X}_i$ includes GPA, entrance-exam percentile, 10th/12th marks, internship sector, and department/degree fixed effects. The coefficient β measures the within-job caste gap in advancement probability at each stage.

Table 1 reports stage-by-stage advancement penalties from Equation 1 with the full control set. In the pooled sample, pre-interview differences are economically small: application screening is -0.25pp ($p < 0.10$), written tests are -0.02pp (n.s.), and group debates are -0.05pp (n.s.). The large break appears only at personal interviews, where the disadvantaged-caste penalty is -29.38pp ($p < 0.01$). In absolute value, the interview-stage coefficient accounts for roughly 99% of the total stage-level penalty mass, so the cumulative success-rate gap is overwhelmingly generated at the final stage.[12]

[12] I operationalize the "interview share" as the absolute value of the interview-to-offer coefficient divided by the sum of absolute values of all four stage-specific coefficients (App→Test, Test→GD, GD→Int, Int→Offer). Each coefficient is the disadvantaged-caste estimate from the stage-specific linear probability model. This mechanical decomposition is what the text summarizes as "over 95%."

Table 1: **Caste Differences in the Success Rate at Each Job Search Stage**

	App → Test (1)	Test → Debate (2)	Debate → Int (3)	Int → Offer (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 1 reports the coefficient β from Equation 1—the effect of belonging to a disadvantaged caste (SC/ST/OBC) on the probability of advancing from one hiring stage to the next, estimated separately for each stage. Each observation is a student-job-year combination for candidates still in contention at that stage. The dependent variable is an indicator for advancing to the next stage. Controls include college GPA, college entrance exam score (JEE rank percentile), 10th and 12th grade marks, and prior internship sector indicators, plus job fixed effects, department fixed effects, degree stream fixed effects, and cohort-year fixed effects. GPA, entrance exam score, and 10th/12th marks enter linearly. When firms conduct common screening rounds for multiple positions, each position is treated as a separate job. Standard errors are clustered at the job level to capture within-job dependence from common screening panels and cutoffs. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

Robustness. The interview penalty persists across both client-facing and non-client-facing roles. Appendix O, Table O.5, Panel A reports a -30.62pp interview penalty ($p < 0.01$) for client-facing jobs (consulting, business development). In the same table, Panel B reports a -28.67pp penalty ($p < 0.01$) for non-client-facing jobs (software engineering, data analytics). The similarity in magnitude across job types is difficult to square with customer discrimination, which would predict materially larger penalties in client-facing roles.

The pattern is also robust to category disaggregation (Appendix B, Table B.1), flexible control specifications—natural splines, decile bins, and saturated department × year × stream fixed effects (Appendix G, Table G.1)—and tests for omitted variable bias using Oster bounds (Appendix H, Table H.1).

Sample Coverage and Selection into Success-Rate Analysis. The success-rate analysis includes 4,005 students (96% of the graduating cohort) who participated in the centralized placement process. The excluded 4% never entered the formal pipeline, so stage-level advancement outcomes are not observed for them.

Appendix O, Table O.6, Panel A evaluates whether this exclusion is likely to inflate the estimated caste penalty. Column (1) shows that the interaction term, Disadvantaged × Did Not Participate in Job Fair, equals -0.044 ($p < 0.05$), indicating that among non-participants,

disadvantaged students have lower GPAs than advantaged students. Because this differential selection implies that omitted disadvantaged students are negatively selected on a key predictor of hiring success, excluding them should bias the disadvantaged-caste coefficient toward zero (less negative), not away from zero. Column (2) then conditions on students who do receive offers and includes job fixed effects. Within the same job, disadvantaged hires have slightly *higher* GPAs (+0.009 log points, $p < 0.01$). The contrast across columns—negative selection among excluded students but positive within-job selection among hired students—is consistent with a higher effective hiring threshold for disadvantaged candidates.

Job Quality. Success rates weight each job equally and therefore recover where screening gaps open, but they do not capture whether those gaps are concentrated in low- or high-paying opportunities. The job-quality analysis adds this economic margin by weighting candidates' opportunity sets by posted salaries: if two groups have similar average advancement rates but diverge most on top-paying jobs, success rates can understate economic loss while job quality will detect it. I therefore measure "job quality" as the average posted salary of jobs each candidate remains in contention for at each stage. Because salaries are *standardized and non-negotiable* within jobs (Section 3.2), variation reflects access rather than negotiation. Stage-specific job-quality regressions are estimated on the set of participants who remain in contention for at least one job at that stage, so sample sizes decline along the pipeline (for example, $N = 4{,}005$ at application, $N = 3{,}031$ at interviews, and $N = 2{,}527$ at offers; Appendix O, Table O.7). Early-stage coefficients therefore describe salary-weighted opportunity-set gaps among students still active at that stage, not only eventual offer recipients.

To estimate the caste gap in job quality *at each stage*, I estimate:

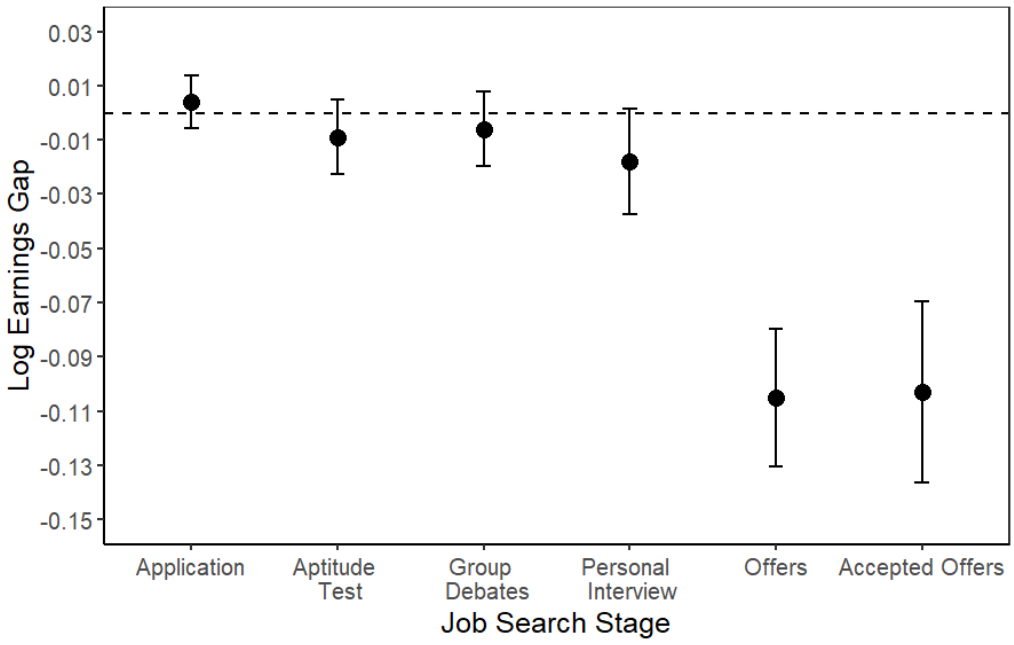
$$\log(\overline{\text{Salary}}_{it}) = \alpha + \beta \cdot \text{Disadvantaged Caste}_i + \mathbf{X}_i'\boldsymbol{\gamma} + \theta_t + \epsilon_{it}, \tag{2}$$

where $\overline{\text{Salary}}_{it}$ is the average posted salary across jobs student i remains eligible for at a given stage in cohort-year t; θ_t are cohort-year fixed effects; and $\mathbf{X}_i$ includes the same controls as Equation 1. Job fixed effects are omitted because the dependent variable is an average across jobs. The coefficient β measures the caste gap in average job quality at each stage.

Figure 2 mirrors the success-rate pattern: gaps are negligible through tests and debates, then become strongly negative at the interview and offer stages, which together account for nearly all of the job-quality gap. Appendix O, Table O.7 reports the full stage-by-stage coefficients. Appendix O, Table O.8 then reports the offer-stage earnings gap (23.1% unconditional; 10.3% with controls), with similar controlled gaps for client- and non-client-facing roles. The magnitude is robust to category disaggregation (Appendix B, Table B.2), flexible controls, and

Oster bounds (Appendix G, Table G.2; Appendix H, Table H.2).

Figure 2: **Job Quality Penalty by Hiring Stage**



Notes: Figure 2 reports the coefficient β from Equation 2 at each hiring stage. Job quality is measured as the average posted salary of jobs a candidate remains in contention for at that stage. Because salaries are *standardized and non-negotiable* within jobs, variation in this average reflects differential job access rather than within-job pay variation. Black dots show the coefficient on disadvantaged caste (SC/ST/OBC indicator), with 95% confidence intervals. Controls include college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, and prior internship sector indicators, plus department fixed effects, degree stream fixed effects, and cohort-year fixed effects. GPA, entrance exam score, and 10th/12th marks enter linearly. Job fixed effects are omitted because the dependent variable is computed at the student-stage level (an average across jobs in contention); salaries are also standardized within jobs, so including job dummies would be both infeasible and redundant.

Sample Selection into Job Quality Analysis. Appendix O, Table O.6, Panel B shows differential selection against disadvantaged candidates on observed ability in the job-quality sample. In column (1), the interaction term Disadvantaged $\times$ Not in Job Quality Regression equals -0.070 ($p < 0.01$), implying that among students excluded from the job-quality regressions (aftermarket or unplaced), disadvantaged students have lower GPAs than excluded advantaged students. This pattern is unlikely to inflate the estimated caste job-quality gap and, if anything, should attenuate it. Column (2) is reported as a reference benchmark and mirrors Panel A: among hired students within the same job, disadvantaged hires have slightly *higher* GPAs (+0.009 log points, $p < 0.01$), consistent with a higher effective hiring threshold.

6.2 Heterogeneity: Penalties are Flat Across Job Selectivity

Section 6.1 documented *average* caste penalties. This section asks whether that interview-stage pattern could instead reflect unobserved *cumulative ability deficits* rather than discriminatory screening, and finds that it does not: in the pooled sample, pre-interview penalties remain near zero across the full selectivity distribution while a large interview penalty appears and stays flat. Section 6.4.2 discusses the key subgroup exception: candidates with caste-identifiable surnames can face earlier penalties.

We observe only advance/not-advance decisions, not stage-level performance scores. Jobs vary widely in selectivity (Appendix O, Table O.9), and more selective jobs draw from the right tail of applicant quality. If cumulative ability deficits were driving the interview result, selectivity should magnify them in two testable ways: (1) pre-interview caste gaps should emerge first in the most selective jobs, and (2) the interview penalty should grow with selectivity. Neither pattern appears. This test rules out broad pre-interview cumulative-deficit stories that should load more strongly in high-cutoff jobs; by itself, it does not rule out interview-specific traits that become observable only in round 4.[13]

To examine heterogeneity across job selectivity, I estimate *at each stage* and *separately for each job j*:

$$\text{Advance}_{ijt} = \alpha_j + \beta_j \cdot \text{Disadvantaged Caste}_i + \mathbf{X}_i'\boldsymbol{\gamma}_j + \theta_t + \epsilon_{ijt}, \tag{3}$$

where Advance_{ijt} is an indicator equal to 1 if student i advances to the next stage for job j in cohort-year t; θ_t are cohort-year fixed effects; and $\mathbf{X}_i$ includes the same individual-level controls as Equation 1: college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, prior internship sector indicators, and department and degree-stream fixed effects (with GPA, entrance exam score, and 10th/12th marks entering linearly). The coefficient β_j is the within-job caste gap in advancement probability at each stage, conditional on observables. I then plot each β_j against the job's cumulative cut rate—the fraction of applicants eliminated by that stage—to trace how the caste penalty changes with selectivity.

Figure 3 presents the results. Two patterns emerge:

1. **No pre-interview penalty, even at highly selective jobs.** At application reading, written tests, and group debates, the caste penalty hovers around zero *across the entire selectivity distribution*. If disadvantaged candidates had lower *cumulative* ability, selective jobs—

[13] This evidence is also hard to reconcile with early-round compensatory favoritism toward disadvantaged candidates: when caste is directly legible from surnames, penalties already appear at application reading (Section 6.4.2), not as early-stage advantages.

which draw from the right tail and thus retain the cumulative best performers—should reveal this through early-stage penalties. They do not.

2. **A large, flat penalty emerges only at interviews.** At the interview stage, disadvantaged candidates face roughly a 30 percentage point penalty that is *nearly constant* across selectivity levels.

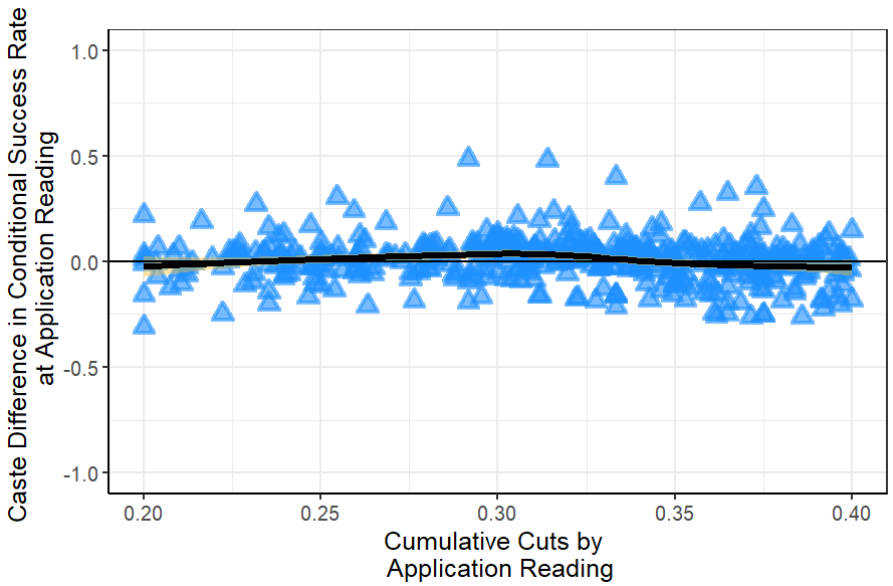
These patterns show low heterogeneity in the caste penalty across the entire distribution of job selectivity, and are inconsistent with *cumulative ability deficits* as an explanation for the interview penalty.

Uncontrolled success rates show the same patterns. Appendix N, Figure N.2 shows *raw* (uncontrolled) success rates—simple caste differences with no controls. At application, a modest raw caste gap exists. As selectivity increases through written tests and group debates, this gap *shrinks toward zero*—then *re-emerges sharply* at interviews, large and flat across selectivity.

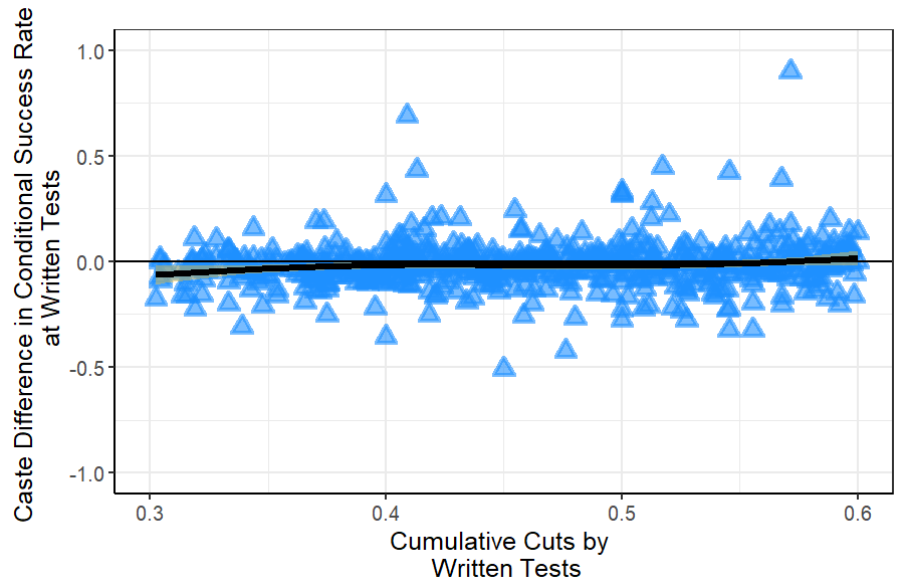
This pattern is particularly compelling because it uses no controls. If disadvantaged candidates had lower cumulative ability, the raw gap should *persist* or *widen* as selectivity increases. Instead, the gap closes across pre-interview stages, implying that disadvantaged candidates who reach interviews are *positively selected*—yet they still face a sharp penalty.

Stylized model shows why flat pre-interview gaps reject cumulative ability deficits. To formalize this selectivity-based logic, I use a stylized threshold model with common pre-interview screening thresholds across groups. In that setup, if disadvantaged candidates had cumulative ability deficits before interviews, pre-interview caste gaps should widen as selectivity rises because higher-cutoff jobs screen deeper into the right tail. The observed pattern—pre-interview gaps that remain near zero across the selectivity distribution—is therefore inconsistent with cumulative-ability-deficit explanations. Appendix I provides the formal setup, propositions, and proofs.

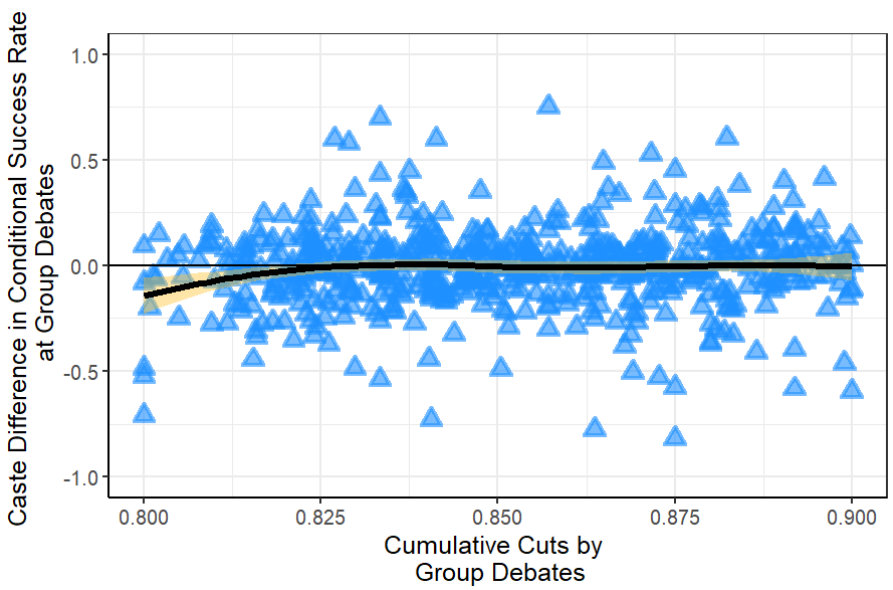
Figure 3: **Caste Penalty in the Conditional Success Rate At Each Screening Round**



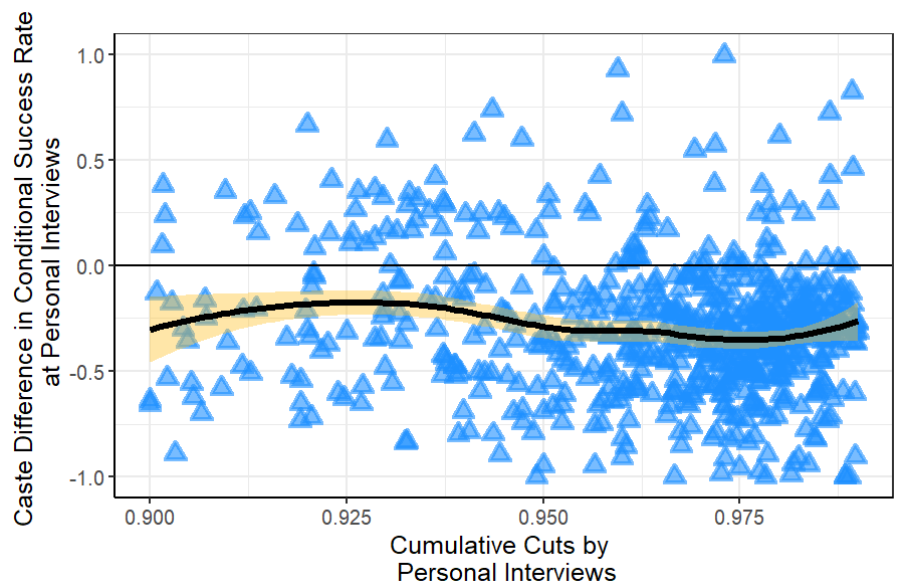
(a) Caste Penalty in Conditional Success Rates versus Cumulative Cuts by Application Reading



(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 3 plots the caste penalty in conditional success rates against job selectivity at each of four screening stages: application reading, written tests, group debates, and personal interviews. Each blue triangle represents a single job. For each job j and stage, I estimate Equation 3: $\text{Advance}_{ijt} = \alpha_j + \beta_j \cdot \text{Disadvantaged Caste}_i + \mathbf{X}_i'\gamma_j + \theta_t + \epsilon_{ijt}$, where controls enter linearly and include college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, prior internship sector, department fixed effects, degree stream fixed effects, and cohort-year fixed effects (for single-year jobs, collinear dummies are automatically dropped). The y-axis shows the estimated β_j for each job; negative values indicate lower advancement rates for disadvantaged castes (SC/ST/OBC) relative to advantaged castes. The x-axis is the stage-specific cumulative cut rate: the cumulative proportion of applicants eliminated by the end of that stage (not the stage-conditional cut). Black lines show smoothed conditional means (LOESS); yellow shading shows 95% confidence intervals. The horizontal line at zero indicates no caste difference. Key finding: the caste penalty hovers around zero at all pre-interview stages regardless of selectivity, but emerges as a large (∼30 percentage point), flat penalty at the interview stage.

6.3 Caste Differences in Promotion and Attrition

Disadvantaged hires have higher first-job promotion rates and longer first-job tenure than advantaged peers, measured over up to five years post-graduation. This horizon is long enough to span typical early-career advancement cycles (entry-level hires at these firms often advance within two to three years), but outcomes are defined within the first employer. The claim is not that interviews reveal no valuable information; rather, the post-hire pattern is difficult to reconcile with disadvantaged hires being weaker on productivity-relevant traits, especially when triangulated with interviews showing recruiters treat retention and advancement as key indicators of hire quality. Lower first-job leaving can partly reflect outside-option or preference differences, but retention is itself a firm-relevant performance margin because turnover imposes replacement and training costs. Accordingly, interpretation does not rely on attrition alone; the promotion evidence and its selectivity pattern move in the same direction.

Measuring Post-Hire Productivity. Career outcomes are measured using promotions and first-job leaving, obtained for over 98% of placed students from alumni records maintained by the placement office and curated from an employment-focused online platform. Appendix E reports coverage benchmarks. These are imperfect productivity proxies, but they are central organizational outcomes in this setting and are commonly used in related evidence on discretionary evaluation (e.g., Hoffman et al., 2018). Turnover is costly—recruiting, training, and disruption all matter for firms—and in entry-level multinational roles with largely standardized pay bands, promotion is a key margin through which firms allocate responsibility and reward realized performance (Bloom et al., 2012; Doeringer and Piore, 1971; Minni, 2024).

Specification. To examine post-hire performance, I estimate:

$$Y_{ijt} = \alpha + \beta \cdot \text{DisadvCaste}_i + \mathbf{X}_i'\boldsymbol{\gamma} + \delta_j + \theta_t + \epsilon_{ijt} \tag{4}$$

where Y_{ijt} is an indicator for promotion (or leaving first job) for individual i in job j and cohort t; δ_j are job fixed effects; θ_t are cohort-year fixed effects; and $\mathbf{X}_i$ includes GPA, entrance exam scores, prior internship sector, and department and degree fixed effects. For promotion outcomes, $\mathbf{X}_i$ also includes job tenure—because caste predicts tenure via differential attrition, the tenure-controlled estimate captures the promotion advantage conditional on equal time at the firm. For the tenure outcome itself, tenure is naturally excluded from controls.

Promotion is coded as 1 if the individual records at least one grade-level increase within five years post-graduation. The post-hire regressions use one observation per graduate's first

placed job; t indexes cohort-year fixed effects rather than person-year panel time. Promotions are identified from title progression within firm and validated using salary changes and benchmark tenure-to-title profiles (Wilmers et al., 2025).

Promotions and First-Job Leaving. Table 2, Panel A, shows that disadvantaged hires have 14–20% higher promotion rates than advantaged peers (columns 1–2).[14] This advantage remains positive—on the order of 15–25%—across job-selectivity tiers (columns 3–9), though estimates become noisier in the most selective bins because sample sizes shrink and not every top-tier bin is individually significant. The key pattern is no systematic attenuation as selectivity rises.[15] At the job-type split (columns 10–11), point estimates remain positive in both client-facing and non-client-facing roles. Results are stable across functional-form checks (Appendix G, Table G.3, Panel A) and Oster bounds (Appendix H, Table H.3, Panel A).

Panel B shows that disadvantaged hires are 17.0 percentage points less likely to leave their first job (28% of the mean). This retention advantage remains negative across selectivity tiers, with wider intervals in the most selective bins, and is directionally similar in client-facing and non-client-facing roles. The estimates remain stable across the same robustness checks (Appendix G, Table G.3, Panel B; Appendix H, Table H.3, Panel B).

Differential Attrition Is Unlikely to Inflate the Estimated Gap. A key concern is whether selective exit among advantaged workers could mechanically inflate the observed promotion advantage. Appendix O, Table O.10 addresses this directly by regressing GPA on leaving status, caste, and their interaction. Three coefficients matter. First, the Disadv. Caste coefficient is +0.007 ($p < 0.01$): within the same job, hired disadvantaged candidates have higher GPAs. Second, Left Job is -0.003 ($p < 0.1$): among advantaged candidates, leavers have slightly lower GPAs than stayers. Third, the interaction (Left Job $\times$ Disadv. Caste = -0.004, n.s.) is near zero: the leaver-stayer GPA gradient does not differ by caste. Together, these estimates indicate that differential attrition is unlikely to generate upward bias in the promotion-gap estimates.

[14] Sample nesting is: 4,005 students in screening-stage success-rate analyses, 2,527 placed students in job-quality analyses, and 2,489 placed students in post-hire promotion/first-job-leaving regressions. The 38-observation reduction from 2,527 to 2,489 reflects placed graduates without usable post-hire career-outcome linkage in the observation window. Coverage and linkage benchmarks are reported in Appendix E; Appendix Table O.10 reports a complementary attrition-selection check within the post-hire sample.

[15] Baseline promotion rates increase with firm selectivity, reflecting stronger internal labor markets and tournament-style advancement at selective firms (Doeringer and Piore, 1971; Lazear and Rosen, 1981).

Interview Triangulation on Quality Metrics. To triangulate interpretation of these post-hire results, I use semi-structured interviews I conducted with senior HR heads overseeing campus recruitment at Tata Consultancy Services (600,000+ employees; operations in 55 countries; LinkedIn), IRIS Software Group (1,000+ employees; 90,000+ customers in 135+ countries; LinkedIn), Tata Motors (10,000+ employees; presence in 125 countries; LinkedIn), Aditya Birla Group (227,500+ employees; operations in 41 countries; LinkedIn), and Jindal South West Steel (40,000+ employees at group level; steel exports to 100+ countries; LinkedIn). This is a limited qualitative sample, so I use it for institutional triangulation rather than stand-alone identification; theme-ordered excerpts appear in Appendix F.

> *"[The campus team] is evaluated on the quality of the overall campus hires and their retention and their cultural fit and their growth in the organization."* (Senior HR head, Aditya Birla Group)

> *"Performance, it is basically the retention part... [w]e invest a lot of time and money on them... if we lose them, then the entire thing is lost."* (Senior HR head, Tata Motors)

> *"[A]fter six months... we do a thorough evaluation and then we decide who are those top performers... [f]or promotion... some will get, some will not get."* (Senior HR head, IRIS Software Group)

> *"[Past] performance rating... feeds into your promotion... [t]he first promotion kicks in minimum three years after."* (Senior HR head, Tata Consultancy Services)

> *"The more important factor that we look at is the attrition from the pool that we have hired... we do look at second year, third year performance as a data point, along with attrition."* (Senior HR head, Jindal South West Steel)

Across these firms, interviewees frame retention, performance progression, and promotion as quality-relevant outcomes rather than automatic post-hire transitions. This aligns with the quantitative evidence here: disadvantaged hires are more likely to be promoted and less likely to leave first jobs across selectivity tiers and job types, which is difficult to reconcile with a pure productivity-deficit account.

Stylized Model: Inframarginality, Latent Productivity, and Higher Interview Thresholds. An important identification concern is inframarginality: even with common interview thresholds, post-hire differences can arise from who survives selection at the margin. Appendix I develops a stylized threshold model (with fuller derivations in Appendix I.1) to make this

logic explicit. The model implies that under equal thresholds, any latent productivity disadvantage should attenuate as job selectivity rises, because both groups' hires are drawn from increasingly similar right tails. By contrast, with a higher effective interview threshold for disadvantaged candidates, stable post-hire advantages (higher promotion and lower first-job leaving / longer tenure) can persist across selectivity tiers because disadvantaged hires are selected from deeper in the right tail. The empirical pattern matches this higher-threshold case, making an equal-threshold latent-productivity account a weaker rationalization of the observed post-hire gaps.

Table 2: **Caste Differences in Promotion and Job Leaving**

	(1) All	(2) All	(3) Top 75%	(4) Top 50%	(5) Top 40%	(6) Top 30%	(7) Top 20%	(8) Top 15%	(9) Top 10%	(10) Client	(11) Non-Client
					By Job Selectivity						*By Job Type*
Panel A: Outcome = Promoted											
Disadv. Caste	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)
Mean Y	0.684	0.684	0.667	0.668	0.676	0.711	0.711	0.706	0.73	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
Panel B: Outcome = Left First Job											
Disadv. Caste	-0.186***	-0.170***	-0.130***	-0.160***	-0.210***	-0.187***	-0.207***	-0.145	-0.132	-0.078*	-0.215***
	(0.025)	(0.027)	(0.029)	(0.037)	(0.043)	(0.055)	(0.079)	(0.092)	(0.121)	(0.045)	(0.034)
Mean Y	0.605	0.605	0.616	0.648	0.655	0.662	0.675	0.668	0.659	0.592	0.611
Observations	2,489	2,489	2,012	1,246	997	748	502	374	252	835	1,654
Adjusted R^2	0.215	0.216	0.24	0.247	0.214	0.179	0.085	0.084	0.054	0.236	0.207
Job FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Other Controls		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Notes: Table 2 reports caste differences in post-hire outcomes for graduates who obtained jobs through the main placement process. Panel A shows promotion rates; Panel B shows rates of leaving the first job. Each observation is one graduate's first placed job, and outcomes are measured over the subsequent five-year window. I estimate: $Y_{ijt} = \alpha + \beta \cdot \text{Disadvantaged Caste}_i + \mathbf{X}_i'\boldsymbol{\gamma} + \delta_j + \theta_t + \epsilon_{ijt}$, where Y_{ijt} is an indicator for promotion (Panel A) or leaving first job (Panel B) for individual i in job j from cohort t; δ_j are job fixed effects; θ_t are cohort-year fixed effects; and $\mathbf{X}_i$ is a vector of controls. Column (1) includes only job fixed effects (δ_j) and cohort-year fixed effects (θ_t). Column (2) adds controls ($\mathbf{X}_i$): college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks (entered linearly), prior internship sector indicators, department fixed effects, degree stream fixed effects, and job tenure (Panel A only; because caste predicts tenure via differential attrition, the tenure-controlled estimate captures the promotion advantage conditional on equal time at the firm rather than the total effect—if anything, omitting tenure would yield a larger advantage). In Panel A, disadvantaged caste graduates are 9.6 percentage points more likely to be promoted (14% of mean). *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

6.4 The Caste Penalty Tracks Cultural-Capital Screening and Caste Revelation

Sections 6.1–6.3 establish that disadvantaged candidates face a higher effective bar at personal interview rounds. The remaining mechanism question is whether this bar operates through caste revelation during interview content, rather than through generic subjectivity. I use two tests. First, I show what personal interviews and group debates actually probe. Second, I use surname timing to test when penalties appear as caste becomes legible.

These tests map directly to the identification target. If personal interviews disproportionately elicit biographical markers of cultural capital, they can make caste legible for the surname-ambiguous majority. If penalties then appear exactly when caste becomes legible—immediately for identifiable surnames and only after personal interview rounds for non-identifiable surnames—the timing supports caste revelation rather than a generic late-stage filter.

6.4.1 LLM Analysis of Mock Interview and Group Debate Transcripts

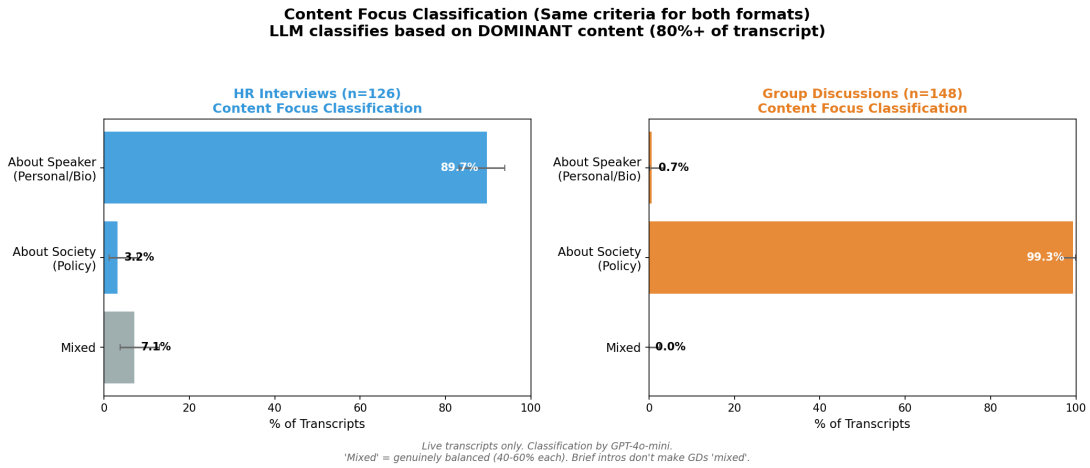
Group debates and personal interview rounds are both subjective and in-person, but only personal interview rounds produce large caste penalties in the hiring data. Because firm-level interview scripts are not observed in the administrative hiring records, I use these transcripts as a content characterization of widely consumed preparation materials rather than as direct transcripts of the focal employers' realized interviews. I analyze 274 live mock transcripts (126 personal interviews, 148 group debates) totaling over 870,000 words (nearly 1 million words), after excluding coaching-only and purely technical sessions.[16] Appendix J documents the full method—prompt design, coding schema, transcript construction, sampling rules, and robustness checks—and details the platform-reach filter: transcripts are drawn from videos with at least 1,000 views, with retained preparation content cumulatively viewed in the hundreds of millions.

Dominant-content level. At this level, each transcript is classified by what occupies at least 80% of the conversation: "about speaker" or "about society." Transcripts that do not cross the 80% threshold in either direction (including 60–80 splits) are coded as "mixed." The coding prompt explicitly distinguishes biography as the topic of exchange from personal

[16] Earlier working drafts used a smaller live subset (183 transcripts: 129 personal interviews and 54 group debates). The current live corpus (274: 126 personal interviews, 148 group debates) reflects expanded group-debate ingestion, stricter live-versus-tips reclassification, and exclusion of personally conducted HR-interview files from the analysis dataset.

anecdotes used rhetorically in policy argumentation. Personal interview rounds are overwhelmingly about the candidate (89.7% coded "about speaker"), whereas group debates are overwhelmingly about policy and society (99.3% coded "about society"). The reverse rates are small (0.7% of group debates coded "about speaker"; 10.3% of personal interviews coded "about society"), with the remainder primarily in the "mixed" category. Figure 4 visualizes this split, and Appendix J, Table J.1 reports the corresponding classification counts. Appendix J, Table J.4 shows that the same high-level split replicates under Anthropic's Claude Sonnet 4 and OpenAI's `gpt-5.2` (runs executed in February 2026). On overlapping transcripts classified as live by each model, dominant-content concordance is high (96.2% GPT-4o-mini vs. Claude Sonnet 4, 98.1% GPT-4o-mini vs. GPT-5.2; Appendix J, Table J.5). Appendix J presents full-size model-specific robustness visuals for Claude Sonnet 4 (Figures J.5–J.7) and GPT-5.2 (Figures J.8–J.10). Those cross-LLM columns use each model's own "live" subset after model-specific live-versus-tips filtering, rather than forcing an identical 274-transcript set. This asymmetry is identification-relevant in the paper's interpretation: the format that systematically extracts biography is also where caste penalties appear.

Figure 4: **Content Focus: Personal Interviews vs. Group Debates**



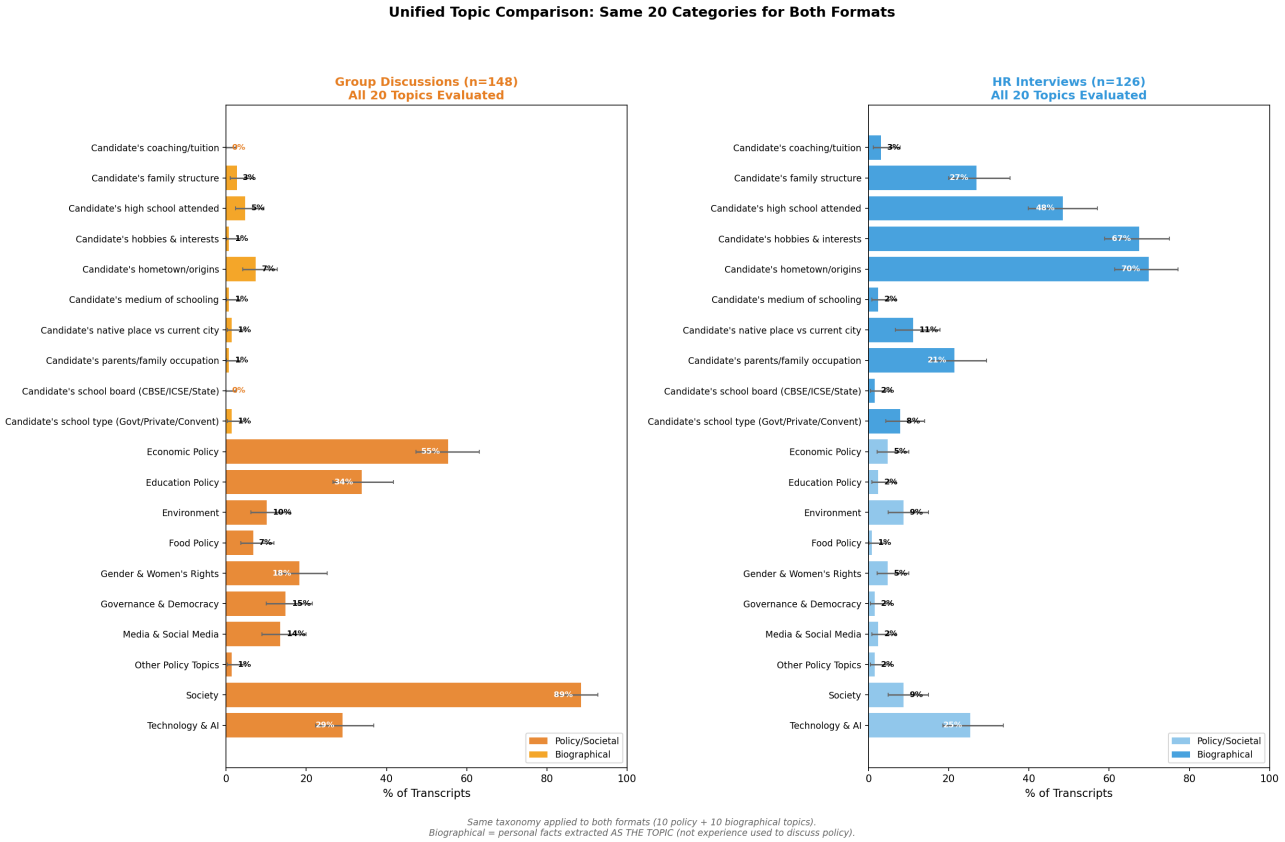
Notes: Figure 4 shows the high-level content focus of 274 live transcripts (126 personal interviews, 148 group debates) from YouTube mock interview preparation videos. Each transcript is classified based on its dominant content (80%+ of transcript): "About Speaker" if the conversation primarily extracts personal/biographical information about the candidate, or "About Society" if it primarily discusses policy/societal topics. Classification was performed using GPT-4o-mini with structured JSON outputs; the prompt explicitly distinguishes biographical content *as the topic* from personal experience used as rhetorical context. Personal interviews overwhelmingly focus on personal/biographical content (89.7%); group debates focus on abstract/policy debates (99.3%). Error bars show 95% Wilson confidence intervals. Results are robust to alternative sampling strategies (paragraph-by-paragraph and random 30% samples) and to cross-LLM validation using Claude Sonnet 4 and `gpt-5.2`. See Appendix J for full methodology, prompts, and robustness figures.

Topic-level evidence under a common taxonomy. Figure 5 applies the same 20-topic taxonomy to both formats and shows the same asymmetry at finer resolution. Personal

interview rounds repeatedly elicit biographical markers linked to cultural capital—hometown, schooling, hobbies, family structure, and parents' occupation—whereas group debates center on policy topics such as economic policy, education policy, governance and democracy, and technology and artificial intelligence. To keep this comparison conservative, biography is coded only when biography is itself the topic of questioning (e.g., "Where are you from?") and not when personal experience is used as context in policy debate. This rule prevents false positives in group debates, where anecdotes often support abstract arguments rather than reveal biography. I also exclude "candidate's college/university education" from the biographical-marker interpretation because all students in this sample have already cleared elite-college admission; that topic appears in 84.9% of personal interview rounds versus 6.8% of group debates. Appendix J reports the full taxonomy, prompt examples, and robustness checks, and Appendix J, Table J.4 confirms that the same HR-versus-GD biographical asymmetry appears under Claude Sonnet 4 and GPT-5.2 (`gpt-5.2`).

Taken together, the dominant-content and topic-level evidence points in the same direction: personal interview rounds are where evaluators systematically observe biographical markers of cultural capital, while group debates are mostly policy-debate settings. This is exactly the mechanism needed for caste revelation among candidates whose surnames are not initially identifiable.

Figure 5: **Topics Probed in Personal Interviews vs. Group Debates**



Notes: Figure 5 shows the percentage of transcripts in which each topic was discussed, using an identical 20-category classification schema applied to both personal interviews and group debates. Sample comprises 274 live transcripts (126 personal interviews, 148 group debates) from YouTube mock interview preparation videos. The y-axis uses a shared alphabetical topic order across both panels (not prevalence ranking). Colors are panel-specific but category-consistent: in the GD panel, dark blue bars are policy/societal topics and orange bars are biographical topics; in the HR panel, light blue bars are policy/societal topics and red bars are biographical topics. Error bars show 95% Wilson confidence intervals. The identical taxonomy ensures comparability: any observed differences reflect genuine content variation rather than differential classification criteria. Classification was performed using GPT-4o-mini with structured JSON outputs; the prompt explicitly distinguishes biographical content *as the topic* (e.g., "Where are you from?") from personal experience used as *context* for policy discussion (e.g., "As someone from a rural area, I think infrastructure investment is crucial"). This distinction prevents false positives where rhetorical moves in group debates are miscoded as biographical extraction. We exclude "candidate's college/university education" from biographical markers because all students in our sample already attend an elite college—this credential was screened at admission, not revealed through interview probing. (For reference, this topic appeared in 84.9% of personal interviews vs. 6.8% of GDs.) See Appendix J for full methodology, prompts, sampling strategies, and cross-LLM robustness checks using Claude Sonnet 4 and `gpt-5.2`.

Limitations, interpretation, and qualitative triangulation. The transcript evidence is indirect: these are mock preparation sessions rather than proprietary firm transcripts, and LLM coding cannot observe all in-room signals (for example posture, confidence display, or tone shifts). Differential access to, or internalization of, coaching resources also cannot be fully ruled out, although that channel alone is difficult to square with the surname-timing pattern and the sharp HR-versus-GD asymmetry under a common coding scheme. The tradeoff is scale and comparability: observing actual interview content at this breadth is rarely feasible,

whereas this corpus provides nearly 1 million words classified under one schema for both formats. This lets us measure content asymmetries systematically rather than from a small set of case interviews.

Interpretation is therefore bounded. If visual/audio cues in group debates were sufficiently unambiguous to reveal caste among surname-ambiguous candidates, penalties should appear there as well; they do not in these data. The same data show that when caste is directly legible from surnames, penalties already appear at application reading. In a setting with weak formal caste-DEI signaling (Appendix C), the absence of group-debate penalties is difficult to reconcile with a story in which generic in-person cues alone already reveal caste at scale.

Interviews I conducted with senior HR heads at Tata Consultancy Services, IRIS Software Group, Tata Motors, Aditya Birla Group, and Jindal South West Steel provide qualitative triangulation that the mock-transcript content is institutionally realistic. Theme-ordered excerpts are reported in Appendix F. Representative statements from each firm point to biographical probing in personal interview rounds:

> *"Definitely in the HR round, we do ask for family backgrounds... we do ask for other things... what basically is the cultural fit of this person."* (Senior HR head, Tata Motors)

> *"I used to check on hobbies... extracurriculars... [and] life skills."* (Senior HR head, Tata Consultancy Services)

> *"[When discussing communication expectations, some candidates] came from backgrounds where their parents were working as domestic help or driver... you can't expect them to read encyclopedias and Enid Blyton."* (Senior HR head, Tata Consultancy Services)

> *"[Do things like hobbies and schooling come up in conversation?] Yes, it does. It does for sure."* (Senior HR head, IRIS Software Group)

> *"These are more aptitude and attitude and aspiration related questions... what's the fit in the organization, that's the basic... [and when asked if interviewers can ask about family background and schooling:] there are no issues on any questions."* (Senior HR head, Aditya Birla Group)

> *"[In personal rounds, interviewers ask about] hobbies, extracurricular[s], [and] what [students] did beyond [the] curriculum... those are more of building rapport [and] trying to understand and gauge... [and] those basically help us take a call."* (Senior HR head, Jindal South West Steel)

This pattern also aligns with evidence from elite U.S. firms, where final-round fit interviews are explicitly framed as biographical and rapport-based screens rather than purely technical assessments:

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

> *"A lot of this job is attitude, not aptitude ... fit is really important ... [you must] like him."* (Rivera, 2012, p. 1008)

Rivera's broader evidence is consistent with this mechanism: in elite U.S. firms, over half of evaluators ranked cultural fit as the most important final-round criterion, and many explicitly used an "airport test" (whether they would want to spend extended informal time with the candidate) (Rivera, 2012). This reinforces the interpretation that final-round personal interviews can combine communication judgments with biographical and rapport-based screening, rather than functioning as purely technical assessments (Lamont, 2009; Turco, 2010).

These qualitative statements are not stand-alone identification, but they align with the transcript evidence that personal interview rounds are where biographical and fit-oriented screening content becomes most salient. Appendix A and Appendix F provide broader legal, scholarly, and qualitative context for this mechanism.

6.4.2 Surname Identifiability Analysis

This second test asks when caste penalties appear as caste becomes legible. If penalties emerge at application reading for caste-identifiable surnames but only after personal interview rounds for non-identifiable surnames, timing identifies caste revelation rather than generic late-stage subjectivity.

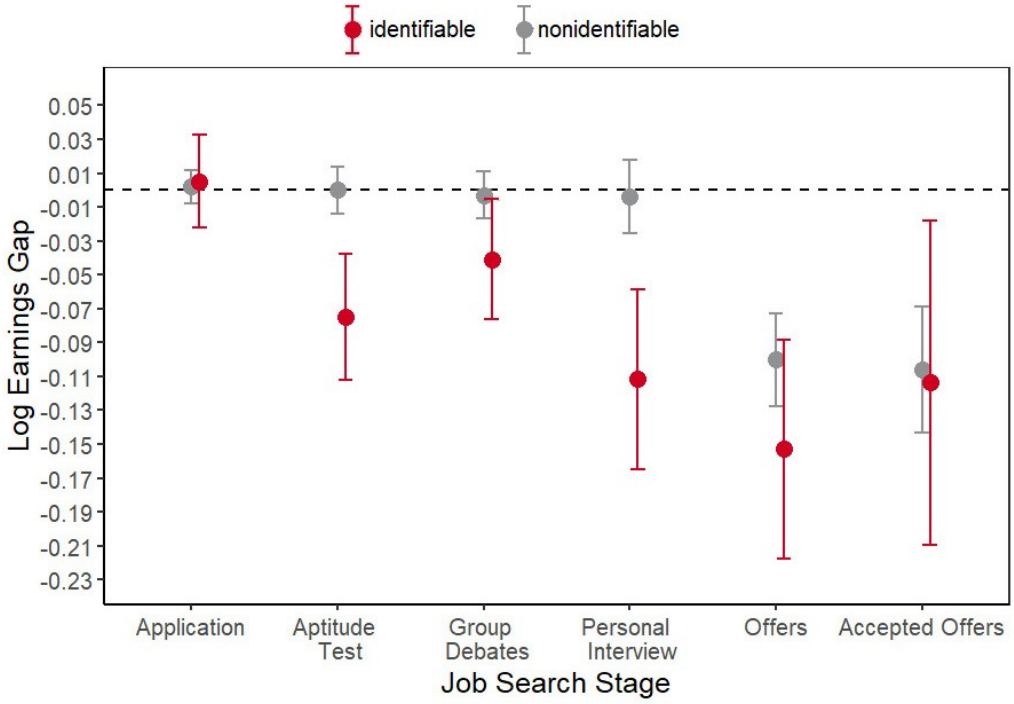
Most surnames are ambiguous in urban settings (e.g., Jodhka, 2015; Parmar, 2020).[17] Under a conservative administrative-list coding, about 15% of students have surnames that unambiguously identify caste status.

How Are Student Names Associated With Caste? We code a surname as *caste-identifiable* only when it provides an unambiguous advantaged/disadvantaged signal based on official administrative lists and corroborating sources. This conservative protocol yields about 15%

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

identifiable surnames and likely understates employers' true ability to infer caste. Appendix K, Table K.1 gives illustrative surname-to-category mappings under this rule, and full source hierarchy and mapping rules are in Appendix K.

Figure 6: **The Caste Job Quality Penalty by Surname Identifiability**



Notes: Figure 6 reports the coefficient β from Equation 2, estimated separately for students with caste-identifiable surnames (red, n=619) and non-identifiable surnames (gray, n=3,386). Each dot shows the caste gap in log average salary of jobs in contention at that stage, with 95% confidence intervals. Since salaries are standardized and non-negotiable within jobs, gaps reflect differences in job access rather than within-job pay variation. All specifications include controls entering linearly: college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, prior internship sector, department fixed effects, degree stream fixed effects, and cohort-year fixed effects. For identifiable surnames, penalties emerge at the written test stage—before any in-person interaction. For non-identifiable surnames, no penalty appears until after personal interviews. Full regression results are in Appendix O, Table O.11, Panels A and B; Oster bounds in Appendix H, Table H.4, Panels A and B; functional form robustness in Appendix G, Table G.4, Panels A and B (success rates), and Appendix G, Table G.5, Panels A and B (job quality).

Robustness of the surname-timing split. Figure 6 shows the job-quality timing split directly: identifiable surnames diverge by the written-test stage, while non-identifiable surnames diverge only after personal interview rounds. Because salaries are standardized, these gaps reflect differences in job access rather than within-job pay. In the success-rate splits (Appendix L, Figures L.1 and L.2), the identifiable-surname penalty appears already at application reading. Appendix O, Table O.11, Panels A and B, reports the stage-wise job-quality split, and Appendix O, Table O.12, Panels A and B, reports the analogous success-rate split.

Selection concern in the identifiable sample. A natural concern is that the identifiable-surname subsample may be differentially selected on baseline productivity, which could mechanically generate earlier penalties. Appendix L, Figures L.1 and L.2 address this directly by comparing raw and controlled application-stage gaps by surname identifiability. In the non-identifiable sample, adding controls for GPA, entrance exam score, school marks, and fixed effects removes the early gap; in the identifiable sample, a sizable early penalty remains after the same controls. Appendix L, Table L.1 then formalizes this selection check: the interaction between disadvantaged caste and non-identifiable surnames in a GPA regression is near zero (0.001), indicating that the caste–GPA slope is not materially different across surname-identifiability groups.

For identifiable surnames, point estimates are often more negative in later rounds, but with uncertainty; the core result is early emergence, not strict monotonic deepening. This does not weaken the revelation mechanism: a name signal at application reading does not imply one-shot deterministic rejection, and later rounds still allow additional evaluators to act on the same information and to combine name salience with newly observed biographical cues. For non-identifiable surnames, by contrast, group debates do not generate actionable caste signals, but personal interview rounds do. Functional-form robustness for surname splits appears in Appendix G, Table G.4, Panels A and B (success rates), and Appendix G, Table G.5, Panels A and B (job quality); Oster robustness is available for the job-quality split in Appendix H, Table H.4, Panels A and B. Together with the transcript evidence, this timing pattern is most consistent with caste revelation through cultural-capital screening. In short personal interviews, these biographical probes are difficult to strategically conceal without creating inconsistencies that interviewers can observe.

What Group Debates Plausibly Do and Do Not Reveal in This Setting. The claim is not that group debates reveal nothing about social background. Accent, fluency, style, and confidence can all carry class-relevant information in interaction (Bourdieu, 1986; Goffman, 1959; Khan, 2011; Lareau, 2003). The empirical claim is narrower: in this setting, whatever group debates reveal is not sufficiently unambiguous and actionable to produce systematic caste penalties for the surname-ambiguous majority. If those signals were already clear enough for consistent penalization, we should observe early-stage penalties analogous to the application-reading penalty observed when caste is directly legible from identifiable surnames. We do not.

This asymmetry also clarifies why strategic concealment is not a complete alternative explanation. Cultural-capital markers are bundled and interactional, not single binary facts that

can be cleanly switched on or off (Bourdieu, 1986; Lamont, 1992; Swidler, 1986). In one-on-one personal interviews, evaluators can probe multiple linked dimensions—schooling trajectory, family background, hometown, hobbies, and communication style—making sustained misrepresentation costly and difficult, especially in socially networked contexts (Jodhka, 2015; Rivera, 2012). These features are consistent with the timing pattern: little actionability in group debates, but clearer caste revelation when biographical probing intensifies in personal interview rounds.

6.5 Discrimination Channels: Statistical, Customer, and Bias

Discrimination channels. Given the interview-stage penalty, the selectivity pattern, and post-hire outcomes, three explanations are central: accurate statistical discrimination, customer discrimination, and bias.

Patterns Less Consistent with Statistical and Customer Discrimination. Customer-discrimination accounts predict stronger penalties in client-facing roles; instead, the interview penalty and post-hire patterns are directionally similar across client-facing and non-client-facing jobs (Sections 6.1 and 6.3). Accurate statistical-discrimination accounts predict attenuation of post-hire gaps as selectivity rises; instead, promotion and retention advantages remain broadly stable across the selectivity distribution (Sections 6.2 and 6.3). The stylized model in Appendix I formalizes this selectivity logic for a broad class of mean- and variance-based screening stories; it narrows those interpretations sharply, while not mechanically ruling out every conceivable higher-moment model without additional structure. Section 6.4 then shows that penalties track caste visibility via surnames and biographical probing in personal interview rounds, rather than generic productivity signals observed at earlier stages.

Higher-Threshold Interpretation and Remaining Ambiguity. Taken together, these patterns are difficult to reconcile with simple customer-discrimination predictions and with a broad class of accurate statistical-discrimination models. They are consistent with a higher effective interview threshold for disadvantaged candidates; that threshold can reflect taste-based discrimination, inaccurate beliefs about group productivity, or both, and the evidence does not mechanically rule out every interview-specific latent trait.

Qualitative Triangulation on Why Bias Can Persist. To triangulate why this bias channel can persist organizationally, I use interviews I conducted with senior HR heads at Tata Consultancy Services, IRIS Software Group, Tata Motors, Aditya Birla Group, and Jindal

South West Steel. These interviews are a limited qualitative sample and not stand-alone identification; I use them for institutional triangulation, with full theme-ordered excerpts in Appendix F. The quotes point to two recurring frictions.

- **Separation between interview screeners and eventual supervisors:**

  "*Predominantly campus hiring will be for a central pool. It is not for any specific business.*" (Senior HR head, IRIS Software Group)

  "*So it's a pool which is hired... which person will go to which location is not very firmed up... it's not the direct manager who is interviewing them.*" (Senior HR head, Tata Motors)

  "*[Graduate engineer trainees] are hired by a pool of panelists... placement in a specific function is mix and match after offers have been made.*" (Senior HR head, Jindal South West Steel)

  "*Those line managers [who interview] are the ones who are available at that time... There is no line manager to student alignment... So it's almost like a lottery.*" (Senior HR head, Aditya Birla Group)

  "*The placements are so random, that it's almost like lottery... there's no interviewer matching done when placement happens.*" (Senior HR head, Tata Consultancy Services)

- **Weak accountability of hiring screeners for downstream performance:**

  "*For talent acquisition [teams], it's mainly the TAT [turnaround time]... [t]heir performance is not linked to the performance of the hired candidates.*" (Senior HR head, Tata Motors)

  "*It is almost impossible to track back their performance, their appraisal or their promotion to the recruitment branch.*" (Senior HR head, Tata Consultancy Services)

  "*After six months of probation we do an evaluation... I do not think too much of data goes to recruitment after six months... beyond that, the ownership will be with the business.*" (Senior HR head, IRIS Software Group)

> *"It's evaluated on the quality of the overall campus hires and their retention and their cultural fit and their growth in the organization."* (Senior HR head, Aditya Birla Group)

> *"There is no qualitative parameters to assess the hiring process itself, except looking at the offer drop ratio."* (Senior HR head, Jindal South West Steel)

These interview accounts do not identify causal effects on their own, but they are consistent with an agency wedge between interview screening and downstream performance feedback, which can allow biased fit assessments to persist.

7 Alternative Explanations and External Validity

This section briefly addresses alternative explanations for the observed patterns and considers the external validity of findings from one elite Indian institution.

Alternative explanations. The main alternatives are hidden cumulative/right-tail ability deficits, accurate statistical discrimination, customer discrimination, strategic delay of discrimination, and selection on unobservables. The evidence pattern is difficult to reconcile with these accounts. If disadvantaged candidates had lower cumulative ability, pre-interview gaps should persist or widen with selectivity; instead pooled pre-interview gaps are near zero even in highly selective jobs (Section 6.2). If accurate statistical discrimination (mean- or variance-based) explained interview-stage penalties, post-hire gaps should attenuate as cutoffs rise; instead promotion and first-job-leaving advantages remain broadly stable across selectivity tiers (Section 6.3; Appendix M). If customer discrimination were primary, penalties should be concentrated in client-facing roles; instead hiring penalties and promotion advantages appear in both client-facing and non-client-facing jobs (Section 6.5).

A complementary channel is reservation-backlash: evaluators may react negatively to public-sector quota discourse even where quotas do not govern these private-sector jobs (Jodhka, 2008; Jodhka and Newman, 2007). But the timing evidence is more consistent with information revelation at interviews than with blanket pre-stage hostility: penalties emerge immediately when surnames reveal caste and otherwise appear primarily when interview content makes caste legible. For additional detail on these alternatives and corresponding checks, see Appendix M.

Institutional context also aligns with this interpretation. Appendix C shows that 81% of sampled recruiting firms (826/1,017) have no substantive DEI page. Among firms with

substantive DEI pages, caste appears in only 5.1% of India-headquartered pages and 1.0% of foreign-headquartered pages (Figure C.1). This formal-policy silence is consistent with weak caste salience in corporate DEI language.

External validity. While the data come from one elite institution, several features support broader relevance. The institution is globally elite in STEM (Martellini et al., 2024), and Section 5 documents that the campus placement architecture is standardized across top Indian colleges. The stage sequence—resume screening, standardized tests, and discretionary final interviews—closely matches elite recruitment pipelines used by multinationals in other settings (Bloom et al., 2012; Rivera, 2015). Evidence from the U.S. and U.K. similarly documents class-linked penalties concentrated at subjective "fit" gates (Laurison and Friedman, 2019; Stansbury and Rodriguez, 2024), and India's civil service shows analogous personality-interview penalties paired with post-hire outperformance (Bhavnani and Lee, 2021). Because this paper studies technical roles with relatively standardized assessments, the estimates are plausibly conservative relative to less standardized hiring environments.

8 Conclusion

Among those hired, disadvantaged-caste graduates exhibit higher promotion rates (roughly 15–20%) and lower first-job leaving, with signs that remain stable across selectivity tiers and across client-facing and non-client-facing roles even where top-tier estimates are less precise. Tracing candidates through the full hiring pipeline at multinational recruiters in India, I find that over 95% of the caste penalty in stage-specific success rates emerges at personal interview rounds, not earlier standardized screens. Lower first-job leaving can partly reflect outside-option differences, but retention remains a firm-relevant margin because turnover is costly; the promotion results point in the same direction. The most coherent interpretation is a higher effective interview threshold for disadvantaged candidates, which implies positive selection among disadvantaged hires.

What do fit interviews assess? The evidence points to cultural-capital screening. Personal interviews systematically extract biographical information—hometown, hobbies, high school attended, family structure, parents' occupation—precisely the markers that Bourdieu (1986) identifies as constitutive of class. At the dominant-content level, personal interviews are overwhelmingly candidate-focused while group debates are overwhelmingly policy-focused; at the topic level, biographical probing is concentrated in personal interviews. This structural difference helps explain why caste penalties emerge specifically at interviews.

Why might this bias persist? Interviews I conducted with senior HR heads at five major multinationals suggest two organizational frictions: interview screeners are commonly separated from eventual supervisors, and talent-acquisition teams are evaluated mainly on process metrics (for example, turnaround time and joins). Some firms monitor cohort-level retention/growth, but interviewers are generally not evaluated on the later performance of their specific hires. These interviews are a limited qualitative sample, but they align with the quantitative evidence on where penalties emerge and why weak feedback can sustain biased fit screening. This is not inevitable. Structured interviews with standardized questions and scoring rubrics can reduce the scope for bias while improving predictive validity (Campion et al., 1997; Schmidt and Hunter, 1998). Linking interviewer assessments to downstream outcomes can also tighten accountability.

Several avenues remain for future work. The heterogeneity analysis (Section 6.2) reveals substantial variation in firm-specific penalties, but which organizational practices—structured interviews, interviewer accountability, supervisor involvement in hiring—predict this variation remains an open question. And while this study focuses on demand-side discrimination conditional on applications, whether disadvantaged candidates also differentially sort across firms is a complementary question.

India's caste system makes visible what remains hidden elsewhere. Class discrimination can operate through similar mechanisms wherever elite hiring relies on subjective "fit" assessments: evaluations of who "belongs" based on biographical markers—accent, credentials, cultural references, family background (Durlauf, 2006; Laurison and Friedman, 2019; Reeves and Friedman, 2024; Rivera, 2015; Stansbury and Rodriguez, 2024). The difference is measurement: caste bundles these capitals into administratively legible categories, exposing patterns that imprecise class proxies obscure. The finding is not that India is unusual, but that India makes legible dynamics likely present wherever subjective fit assessments occur. "Every established order tends to produce ... the naturalization of its own arbitrariness" (Bourdieu, 1977, p. 164). Elite hiring's language of "fit" is one such naturalization—cultural-capital screening that often converts inherited status into signals interpreted as merit, sustaining bias under the language of meritocracy.

References

Altonji, J. G. and R. M. Blank (1999). Race and gender in the labor market. *Handbook of Labor Economics 3*, 3143–3259.

Asher, S., P. Novosad, and C. Rafkin (2024). Intergenerational mobility in india: New

measures and estimates across time and social groups. *American Economic Journal: Applied Economics 16*(2), 1–25.

Ayres, I. (2002). Outcome Tests of Racial Disparities in Police Practices. *Justice Research and Policy 4*(1-2), 131–142.

Bagde, S., D. Epple, and L. Taylor (2016). Does affirmative action work? caste, gender, college quality, and academic success in india. *American Economic Review 106*(6), 1495–1521.

Banerjee, A., E. Duflo, M. Ghatak, and J. Lafortune (2013). Marry for what? caste and mate selection in modern india. *American Economic Journal: Microeconomics 5*(2), 33–72.

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.

Basham, A. L. (1954). *The Wonder That Was India: A Survey of the Culture of the Indian Sub-Continent Before the Coming of the Muslims*. London: Sidgwick and Jackson. Classic survey of pre-Islamic Indian civilization.

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

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

Bertrand, M. and E. Duflo (2017). Field experiments on discrimination. In A. V. Banerjee and E. Duflo (Eds.), *Handbook of Economic Field Experiments*, Volume 1, pp. 309–393. Amsterdam: North-Holland.

Bertrand, M., R. Hanna, and S. Mullainathan (2010). Affirmative action in education: Evidence from engineering college admissions in india. *Journal of Public Economics 94*(1-2), 16–29.

Bertrand, M. and S. Mullainathan (2004). Are Emily and Greg more employable than Lakisha and Jamal? a field experiment on labor market discrimination. *American Economic Review 94*(4), 991–1013.

Béteille, A. (1965). *Caste, Class, and Power: Changing Patterns of Stratification in a Tanjore Village*. Berkeley: University of California Press.

Béteille, A. (1996). Caste in contemporary india. *Caste Today*, 150–179. Argues jati corresponds to ethnicity rather than caste.

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.

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.

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

Blau, P. M. and O. D. Duncan (1967). *The American Occupational Structure*. New York: Wiley.

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

Bourdieu, P. (1977). *Outline of a Theory of Practice*. Cambridge: Cambridge University Press.

Bourdieu, P. (1984). *Distinction: A Social Critique of the Judgement of Taste*. Cambridge, MA: Harvard 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. New York: Greenwood.

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

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.
Cassan, G., D. Keniston, and T. Kleineberg (2021). A division of laborers: Identity and efficiency in india. Working Paper 28462, National Bureau of Economic Research. R&R at Econometrica.
Chatman, J. A. (1991). Matching people and organizations: Selection and socialization in public accounting firms. *Administrative Science Quarterly 36*(3), 459–484.
Clark, G. (2014). *The Son Also Rises: Surnames and the History of Social Mobility*. Princeton, NJ: Princeton University Press.
Cohn, B. S. (1987). The census, social structure and objectification in South Asia. In *An Anthropologist Among the Historians and Other Essays*, pp. 224–254. Delhi: Oxford University Press.
Datta, S. (2017). *A History of the Indian University System: Emerging from the Shadows of the Past*. Palgrave Macmillan.
Deshpande, A. (2011a). *The Grammar of Caste: Economic Discrimination in Contemporary India*. New Delhi: Oxford University Press.
Deshpande, A. (2011b). *The Grammar of Caste: Economic Discrimination in Contemporary India*. New Delhi: Oxford University Press.
Dirks, N. B. (2001). *Castes of Mind: Colonialism and the Making of Modern India*. Princeton, NJ: Princeton University Press.
Doeringer, P. B. and M. J. Piore (1971). *Internal Labor Markets and Manpower Analysis*. Lexington, MA: Heath Lexington Books.
Dumont, L. (1966). *Homo Hierarchicus: The Caste System and Its Implications*. Chicago: University of Chicago Press. English translation 1970.
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.
Erikson, R. and J. H. Goldthorpe (1992). *The Constant Flux: A Study of Class Mobility in Industrial Societies*. Oxford: Oxford University Press.
Fitzgerald, F. S. (1925). *The Great Gatsby*. New York: Charles Scribner's Sons.
Galanter, M. (1984). *Competing Equalities: Law and the Backward Classes in India*. Berkeley: University of California Press.
Ghurye, G. S. (1969). *Caste and Race in India* (5th ed.). Bombay: Popular Prakashan.
Goffman, E. (1959). *The Presentation of Self in Everyday Life*. Garden City, NY: Doubleday.
Goldberg, A., S. B. Srivastava, V. G. Manian, W. Monroe, and C. Potts (2016). Fitting in or standing out? The tradeoffs of structural and cultural embeddedness. *American Sociological Review 81*(6), 1190–1222.
Government of India (1950). The constitution of india (as amended). Relevant provisions: Articles 15(4), 16(4), 340, 341, 342, and 338B.
Highhouse, S. (2008). Stubborn reliance on intuition and subjectivity in employee selection. *Industrial and Organizational Psychology 1*(3), 333–342.
His Majesty in Council (1936). The government of india (scheduled castes) order, 1936. Order in Council under the Government of India Act, 1935.
Hoffman, M., L. B. Kahn, and D. Li (2018). Discretion in Hiring. *The Quarterly Journal of Economics 133*(2), 765–800.
Hoffman, M. and C. T. Stanton (2024). People, practices, and productivity: A review of new advances in personnel economics. NBER Working Paper 32849.

Inden, R. (2001). *Imagining India* (2nd ed.). Bloomington: Indiana University Press. Notes no universally accepted definition of caste.

Jodhka, S. S. (2008, October). Caste & the corporate sector. *Indian Journal of Industrial Relations 44*(2), 185–193. Invited article; JSTOR stable/27768188.

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.

Karabel, J. (2005). *The Chosen: The Hidden History of Admission and Exclusion at Harvard, Yale, and Princeton*. Boston: Houghton Mifflin.

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

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.

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.

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.

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

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

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

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

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

Mamgain, R. P. (2019). Growth, poverty and non-farm employment in india. *Indian Journal of Labour Economics 62*(2), 197–214.

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

Mandal Commission (1980). The Mandal Commission Report. Prime Minister's High Level Committee, Cabinet Secretariat, Government of India.

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

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.

Miller, L. C. (1919). Report of the committee appointed to consider steps necessary for the adequate representation of communities in the public service.

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

Munshi, K. (2019). Caste and the indian economy. *Journal of Economic Literature 57*(4), 781–834.

Omi, M. and H. Winant (2014). *Racial Formation in the United States* (3rd ed.). New York: Routledge.

O'Reilly, C. A., J. Chatman, and D. F. Caldwell (1991). People and organizational culture: A profile comparison approach to assessing person–organization fit. *Academy of Management Journal 34*(3), 487–516.

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

Parkin, F. (1979). *Marxism and Class Theory: A Bourgeois Critique*. London: Tavistock.

Parmar, D. (2020). Caste and the politics of perception in india's labor markets. *World Development 135*, 105066.

Pew Research Center (2021). Religion in india: Tolerance and segregation. Technical report, Pew Research Center.

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.

Rajya Sabha Committee on Personnel, Public Grievances, Law and Justice (2005). Eighth report on the scheduled castes, scheduled tribes and other backward classes (reservation in posts and services) bill, 2004. Technical report, Rajya Sabha. Historical background on Mysore's 1918–21 policy and the Miller Committee.

Reeves, A. and S. Friedman (2024). *Born to Rule: The Making and Remaking of the British Elite*. Cambridge, MA: Belknap Press of Harvard University Press.

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

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.

Sackett, P. R., C. Zhang, C. M. Berry, and F. Lievens (2022). Revisiting meta-analytic estimates of validity in personnel selection: Addressing systematic overcorrection for restriction of range. *Journal of Applied Psychology 107*(12), 2040–2068.

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

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

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.

Schmidt, F. L. and J. E. Hunter (1998). The validity and utility of selection methods in personnel psychology: Practical and theoretical implications of 85 years of research findings. *Psychological Bulletin 124*(2), 262–274.

Small, M. L. and D. Pager (2020). Sociological perspectives on racial discrimination. *Journal of Economic Perspectives 34*(2), 49–67.

Srinivas, M. N. (1966). *Social Change in Modern India*. Berkeley: University of California Press.

Srivastava, S. B., A. Goldberg, V. G. Manian, and C. Potts (2018). Enculturation trajectories: Language, cultural adaptation, and individual outcomes in organizations. *Management Science 64*(3), 1348–1364.

Stansbury, A. (2024). Social class and career progression in elite professions. *American Economic Journal: Applied Economics Forthcoming*.

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

Stevens, M. P. (2009). *Creating a Class: College Admissions and the Education of Elites*. Harvard University Press.

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

Supreme Court of India (1992). Indra sawhney v. union of india. AIR 1993 SC 477.

Swartz, D. L. (1997). *Culture and Power: The Sociology of Pierre Bourdieu*. Chicago: University of Chicago Press.

Swidler, A. (1986). Culture in action: Symbols and strategies. *American Sociological Review 51*(2), 273–286.

Tambe, P. and L. M. Hitt (2012). The productivity of information technology investments: New evidence from IT labor data. *Information Systems Research 23*(3), 599–617.

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

Thorat, S. and K. S. Newman (Eds.) (2010). *Blocked by Caste: Economic Discrimination in Modern India*. New Delhi: Oxford University Press.

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

Weber, M. (1978a). *Economy and Society: An Outline of Interpretive Sociology*. Berkeley: University of California Press.

Weber, M. (1978b). *Economy and Society: An Outline of Interpretive Sociology*. Berkeley: University of California Press.

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.

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

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

Appendix

Table of Contents

A Legal and Scholarly Foundations on Caste and Class **51**

A.1 Legal Architecture of Administrative Coding 51

A.2 Scholarly Foundations for a Class-Relevant Mechanism 54

A.3 Comparative Implications for This Paper 56

B Robustness to Category Aggregation **58**

C Caste in Corporate DEI Policy **60**

D Representativeness of the Job Fair at the Elite College **63**

D.1 Systematic Evidence of Standardization Across Elite Institutions 63

D.2 Corroborating Evidence 63

E Employment-Platform Coverage Benchmarks and Profile-Matching Methodology **65**

E.1 Career Outcome Data Collection and Professional Profile Matching Methodology 65

F HR Professional Interviews: Theme-Ordered Excerpts **66**

G Robustness to Functional Form **69**

G.1 Success Rates 69

G.2 Job Quality 70

G.3 Post-Hire Outcomes 71

G.4 Success Rates by Surname Identifiability 72

G.5 Job Quality by Surname Identifiability 73

H Robustness to Omitted Variable Bias: Oster Bounds **75**

H.1 Success Rates 75

H.2 Job Quality 76

H.3 Promotion and Tenure Rates 77

H.4 Job Quality by Surname Identifiability 79

I A Model Exploiting Plausibly Exogenous Variation in Selectivity Across Jobs to Identify Discrimination **80**

I.1 Extending the Simple Model of Late-Stage Discrimination with Imperfect Information 86

I.2 Derivation of Expected Performance Above Threshold 89

J LLM-Based Transcript Analysis Methodology **91**

J.1 Key Design Principle: Unified Taxonomy 91

J.2 Data Collection 91

J.3 Classification Prompts 92

J.4 Classification Process 95

J.5 Results Summary 95

J.6 Sampling-Based Sensitivity Analysis 97

K Detailed Methodology for Classifying Caste-Identifiable Surnames **111**

K.1 Hierarchy of Data Sources 111

K.2 Administrative Framework and the Definition of Disadvantage 112

K.3 Operational Mapping and Conservative Criteria 112

L Differential Selection Across Surname Samples **115**

M More Details on Alternative Explanations **119**

N Additional Figures **120**

O Additional Tables **122**

A Legal and Scholarly Foundations on Caste and Class

Discrimination research on caste, class, and race often proceeds in parallel traditions with limited cross-walk: South Asian scholarship develops the historical and institutional specifics of caste; stratification sociology develops the language of class and cultural capital; race scholarship emphasizes visual classification and group boundaries. This appendix integrates those traditions for a narrow empirical purpose: for the interview-stage mechanism studied here, India's legal-administrative coding and comparative scholarship justify treating caste as a tractable measure of cumulative class-relevant disadvantage.

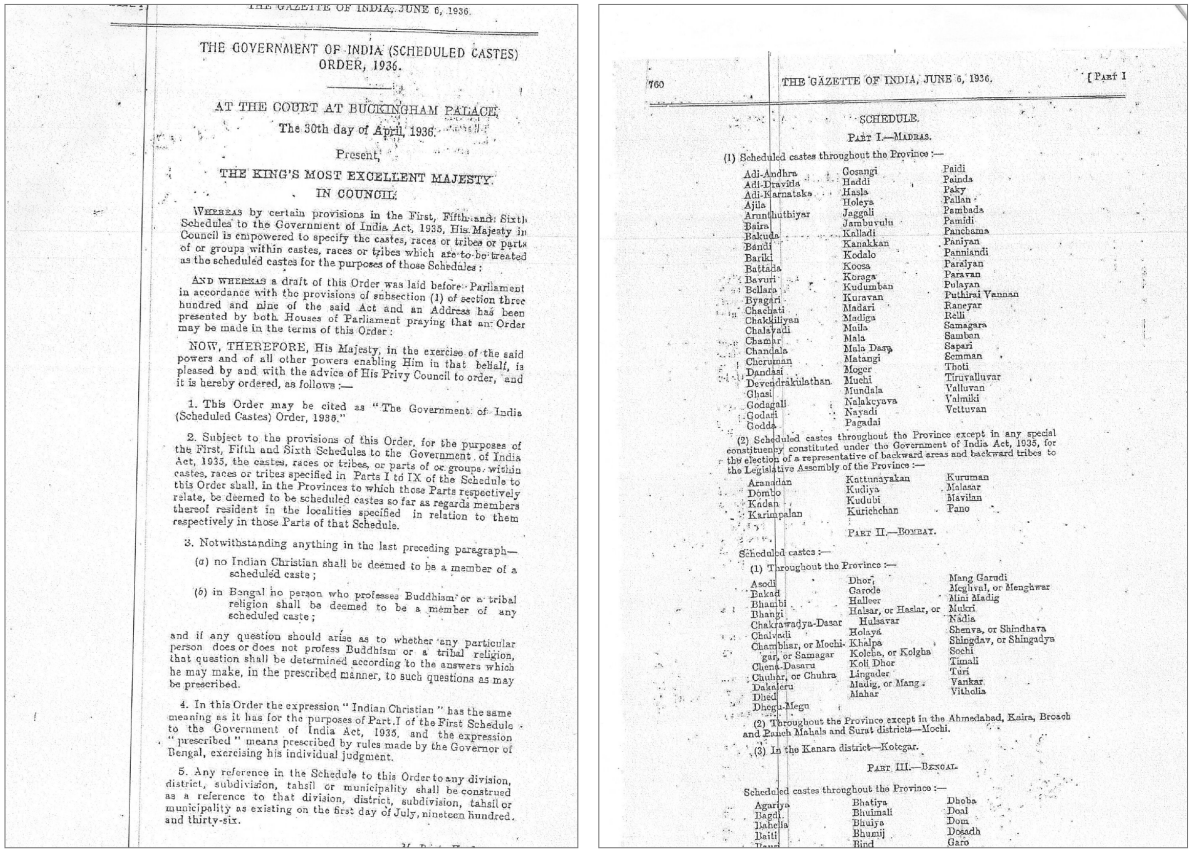
A.1 Legal Architecture of Administrative Coding

Pre-independence bridge: "Depressed Classes" to "Scheduled Castes." Before constitutional codification, reform-era and colonial administrative documents used the term "Depressed Classes" to describe communities subjected to untouchability and associated civic exclusion. Early representation politics in Mysore—including memorials submitted to Maharaja Krishnaraja Wadiyar IV and the subsequent Miller Committee process chaired by Sir Leslie C. Miller—already treated group disadvantage as an administrative governance problem rather than as a purely ritual abstraction (Miller, 1919; Rajya Sabha Committee on Personnel, Public Grievances, Law and Justice, 2005). At the all-India level, the shift to schedule-based language was formalized in the Government of India (Scheduled Castes) Order, 1936 under the Government of India Act, 1935 (His Majesty in Council, 1936). This is the institutional origin of the term "Scheduled Castes": communities listed in statutory schedules.

Figure A.1: **Historical Schedule Evidence: Government of India (Scheduled Castes) Order, 1936**

Official Archive: Government of India (Scheduled Castes) Order, 1936

Left: order title page. Right: schedule listing by province.



THE GAZETTE OF INDIA, JUNE 6, 1936.

THE GOVERNMENT OF INDIA (SCHEDULED CASTES) ORDER, 1936.

AT THE COURT AT BUCKINGHAM PALACE, The 30th day of April, 1936.

Present,

THE KING'S MOST EXCELLENT MAJESTY IN COUNCIL

WHEREAS by certain provisions in the First, Fifth and Sixth Schedules to the Government of India Act, 1935, His Majesty in Council is empowered to specify the castes, races or tribes or parts of or groups within castes, races or tribes which are to be treated as the scheduled castes for the purposes of those Schedules:

AND WHEREAS a draft of this Order was laid before Parliament in accordance with the provisions of subsection (1) of section three hundred and nine of the said Act and an Address has been presented by both Houses of Parliament praying that an Order may be made in the terms of this Order:

NOW, THEREFORE, His Majesty, in the exercise of the said powers and of all other powers enabling Him in that behalf, is pleased by and with the advice of His Privy Council to order, and it is hereby ordered, as follows:—

1. This Order may be cited as "The Government of India (Scheduled Castes) Order, 1936."

2. Subject to the provisions of this Order, for the purposes of the First, Fifth and Sixth Schedules to the Government of India Act, 1935, the castes, races or tribes, or parts of or groups within castes, races or tribes specified in Parts I to IX of the Schedule to this Order shall, in the Provinces to which those Parts respectively relate, be deemed to be scheduled castes so far as regards members thereof resident in the localities specified in relation to them respectively in those Parts of that Schedule.

3. Notwithstanding anything in the last preceding paragraph—

(a) no Indian Christian shall be deemed to be a member of a scheduled caste;

(b) in Bengal no person who professes Buddhism or a tribal religion shall be deemed to be a member of any scheduled caste;

and if any question should arise as to whether any particular person does or does not profess Buddhism or a tribal religion, that question shall be determined according to the answers which he may make, in the prescribed manner, to such questions as may be prescribed.

4. In this Order the expression "Indian Christian" has the same meaning as it has for the purposes of Part I of the First Schedule to the Government of India Act, 1935, and the expression "prescribed" means prescribed by rules made by the Governor of Bengal, exercising his individual judgment.

5. Any reference in the Schedule to this Order to any division, district, subdivision, tahsil or municipality shall be construed as a reference to that division, district, subdivision, tahsil or municipality as existing on the first day of July, nineteen hundred and thirty-six.

760 THE GAZETTE OF INDIA, JUNE 6, 1936. [PART I

SCHEDULE.

PART I.—MADRAS.

(1) Scheduled castes throughout the Province:—

Adi-Andhra	Gosangi	Paidi
Adi-Dravida	Haddi	Painda
Adi-Karnataka	Hasla	Paky
Ajila	Holeya	Pallan
Arunthuthiyar	Jaggali	Pambada
Baira	Jambuvulu	Pamidi
Bakuda	Kalladi	Panchama
Bandi	Kanakkan	Paniyan
Barike	Kodalo	Panniandi
Battada	Koosa	Paraiyan
Bavuri	Koraga	Paravan
Bellara	Kudumban	Pulayan
Byagari	Kuravan	Puthirai Vannan
Chachati	Madari	Raneyar
Chakkiliyan	Madiga	Relli
Chalvadi	Malla	Samagara
Chambr	Mala	Samban
Chandala	Mala Dasa	Sapari
Cheruman	Matangi	Semman
Dandasi	Moger	Thoti
Devendrakulathan	Muchi	Tiruvalluvar
Ghasi	Mundala	Valluvan
Godagali	Nalakeyava	Valmiki
Godari	Nayadi	Vettuvan
Godda	Pagadai	

(2) Scheduled castes throughout the Province except in any special constituency constituted under the Government of India Act, 1935, for the election of a representative of backward areas and backward tribes to the Legislative Assembly of the Province:—

Aranadan	Kattunayakan	Kuruman
Domlo	Kudiya	Malasar
Kadan	Kudubi	Mavilan
Karimpalan	Kurichchan	Pano

PART II.—BOMBAY.

Scheduled castes:—

(1) Throughout the Province:—

Asodi	Dhor	Mang Garudi
Bakadi	Garode	Mcghwal, or Menghwar
Bhadbi	Halleer	Mini Madig
Bhangi	Halsar, or Haslar, or Hulsavar	Mukri
Chakravadya-Dasar	Holeya	Nadia
Chalvadi	Khalpa	Shenva, or Shindhava
Chambhar, or Mochigar, or Samagar	Kolcha, or Kolgha	Shingdav, or Shingadya
Chena-Dasaru	Koli Dhor	Sochi
Chuhar, or Chuhra	Lingadar	Timali
Dakaleru	Madig, or Mang	Turi
Dhed	Mahar	Vankar
Dhegu-Megu		Vithalia

(2) Throughout the Province except in the Ahmedabad, Kaira, Broach and Panch Mahals and Surat districts—Mochi.

(3) In the Kanara district—Kotegar.

PART III.—BENGAL.

Scheduled castes throughout the Province:—

Agariya	Bhatiya	Dhoba
Bagdi	Bhumali	Doai
Bahelia	Bhuiya	Dom
Baiti	Bhumij	Dosadh
Bauri	Bind	Garo

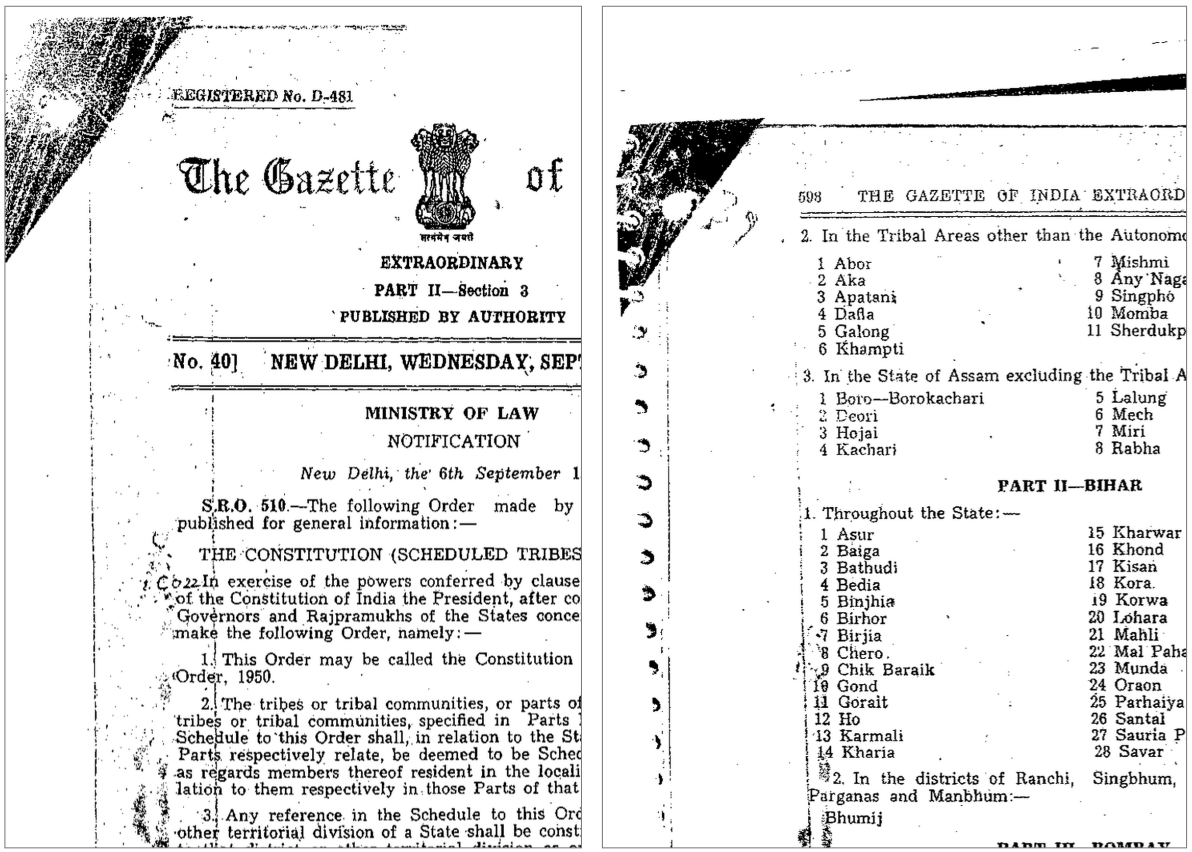
Notes: Reproduction from official government archival material for the Government of India (Scheduled Castes) Order, 1936. Left panel shows the order title page; right panel shows schedule listings by province. This is documentary evidence for the literal schedule-based origin of the term "Scheduled Castes." Source: Ministry of Social Justice and Empowerment archival copy (https://socialjustice.gov.in/public/ckeditor/upload/82951673327147.PDF), accessed February 19, 2026.

Constitutional continuation and post-independence categories. Post-independence law retained and expanded this coding architecture. Articles 341 and 342 govern Scheduled Castes and Scheduled Tribes through Presidential Orders, while Article 340 empowers commissions on socially and educationally backward classes (Government of India, 1950). In other words, India's constitutional design explicitly recognizes that disadvantage is grouped and administered through legally legible social categories. This does not erase variation within categories; it creates a policy-relevant administrative frame that courts, commissions, and public institutions repeatedly use (Galanter, 1984; Bayly, 2008; Dirks, 2001; Cohn, 1987).

Figure A.2: **Historical Schedule Evidence: Constitution (Scheduled Tribes) Order, 1950**

Official Archive: Constitution (Scheduled Tribes) Order, 1950

Left: order title page. Right: schedule listing by state/province.



REGISTERED No. D-481

The Gazette of

EXTRAORDINARY

PART II—Section 3

PUBLISHED BY AUTHORITY

No. 40] NEW DELHI, WEDNESDAY, SEP

MINISTRY OF LAW

NOTIFICATION

New Delhi, the 6th September 1

S.R.O. 510.—The following Order made by published for general information:—

THE CONSTITUTION (SCHEDULED TRIBES

In exercise of the powers conferred by clause of the Constitution of India the President, after co Governors and Rajpramukhs of the States conce make the following Order, namely:—

1. This Order may be called the Constitution Order, 1950.

2. The tribes or tribal communities, or parts of tribes or tribal communities, specified in Parts Schedule to this Order shall, in relation to the St Parts respectively relate, be deemed to be Schec as regards members thereof resident in the locali lation to them respectively in those Parts of that

3. Any reference in the Schedule to this Orc other territorial division of a State shall be const

598 THE GAZETTE OF INDIA EXTRAORD

2. In the Tribal Areas other than the Autonom

1 Abor
2 Aka
3 Apatani
4 Dafla
5 Galong
6 Khampti
7 Mishmi
8 Any Naga
9 Singpho
10 Momba
11 Sherdukp

3. In the State of Assam excluding the Tribal A

1 Boro—Borokachari
2 Deori
3 Hojai
4 Kachari
5 Lalung
6 Mech
7 Miri
8 Rabha

PART II—BIHAR

1. Throughout the State:—

1 Asur
2 Baiga
3 Bathudi
4 Bedia
5 Binjhia
6 Birhor
7 Birjia
8 Chero
9 Chik Baraik
10 Gond
11 Gorait
12 Ho
13 Karmali
14 Kharia
15 Kharwar
16 Khond
17 Kisan
18 Kora
19 Korwa
20 Lohara
21 Mahli
22 Mal Paha
23 Munda
24 Oraon
25 Parhaiya
26 Santal
27 Sauria P
28 Savar

2. In the districts of Ranchi, Singbhum, Parganas and Manbhum:—

Bhumij

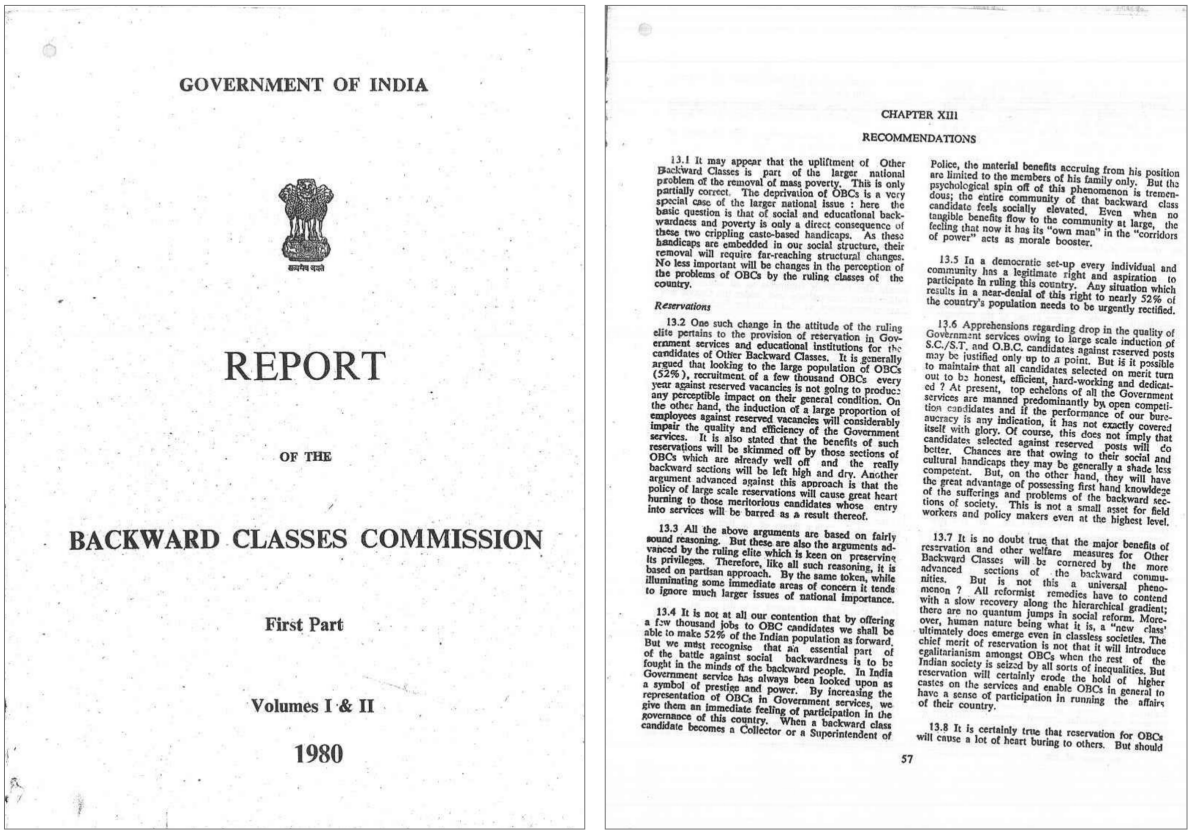
Notes: Reproduction from official government archival material for the Constitution (Scheduled Tribes) Order, 1950. Left panel shows the order title page; right panel shows schedule listings by state/province. This provides documentary evidence for the same schedule-based legal coding logic for ST categories. Source: Ministry of Tribal Affairs official copy (http://www.tribal.nic.in/downloads/CLM/CLM_1/1.pdf), accessed February 19, 2026.

Why Mandal matters for measurement in this paper. The Mandal Commission is central because it operationalized OBC identification through an eleven-indicator composite spanning social, educational, and economic domains, with stronger weight on social and educational backwardness than on income alone (Mandal Commission, 1980; Galanter, 1984). That weighting is exactly why Mandal is empirically relevant here: it approximates cumulative social position rather than a single contemporaneous income measure. For an interview-screening paper, this matters because employers often infer class through long-horizon signals (schooling trajectory, speech, networks, family background), not just current resources.

Figure A.3: **Historical OBC Evidence: Backward Classes Commission (Mandal) Report, 1980**

Official Archive: Backward Classes Commission (Mandal) Report, 1980

Left: report cover page. Right: recommendations chapter page.



GOVERNMENT OF INDIA

REPORT

OF THE

BACKWARD CLASSES COMMISSION

First Part

Volumes I & II

1980

CHAPTER XIII

RECOMMENDATIONS

13.1 It may appear that the upliftment of Other Backward Classes is part of the larger national problem of the removal of mass poverty. This is only partially correct. The deprivation of OBCs is a very special case of the larger national issue : here the basic question is that of social and educational backwardness and poverty is only a direct consequence of these two crippling caste-based handicaps. As these handicaps are embedded in our social structure, their removal will require far-reaching structural changes. No less important will be changes in the perception of the problems of OBCs by the ruling classes of the country.

Reservations

13.2 One such change in the attitude of the ruling elite pertains to the provision of reservation in Government services and educational institutions for the candidates of Other Backward Classes. It is generally argued that looking to the large population of OBCs (52%), recruitment of a few thousand OBCs every year against reserved vacancies is not going to produce any perceptible impact on their general condition. On the other hand, the induction of a large proportion of employees against reserved vacancies will considerably impair the quality and efficiency of the Government services. It is also stated that the benefits of such reservations will be skimmed off by those sections of OBCs which are already well off and the really backward sections will be left high and dry. Another argument advanced against this approach is that the policy of large scale reservations will cause great heart burning to those meritorious candidates whose entry into services will be barred as a result thereof.

13.3 All the above arguments are based on fairly sound reasoning. But these are also the arguments advanced by the ruling elite which is keen on preserving its privileges. Therefore, like all such reasoning, it is based on partisan approach. By the same token, while illuminating some immediate areas of concern it tends to ignore much larger issues of national importance.

13.4 It is not at all our contention that by offering a few thousand jobs to OBC candidates we shall be able to make 52% of the Indian population as forward. But we must recognise that an essential part of of the battle against social backwardness is to be fought in the minds of the backward people. In India Government service has always been looked upon as a symbol of prestige and power. By increasing the representation of OBCs in Government services, we give them an immediate feeling of participation in the governance of this country. When a backward class candidate becomes a Collector or a Superintendent of Police, the material benefits accruing from his position are limited to the members of his family only. But the psychological spin off of this phenomenon is tremendous; the entire community of that backward class candidate feels socially elevated. Even when no tangible benefits flow to the community at large, the feeling that now it has its "own man" in the "corridors of power" acts as morale booster.

13.5 In a democratic set-up every individual and community has a legitimate right and aspiration to participate in ruling this country. Any situation which results in a near-denial of this right to nearly 52% of the country's population needs to be urgently rectified.

13.6 Apprehensions regarding drop in the quality of Government services owing to large scale induction of S.C./S.T. and O.B.C. candidates against reserved posts may be justified only up to a point. But is it possible to maintain that all candidates selected on merit turn out to be honest, efficient, hard-working and dedicated ? At present, top echelons of all the Government services are manned predominantly by open competition candidates and if the performance of our bureaucracy is any indication, it has not exactly covered itself with glory. Of course, this does not imply that candidates selected against reserved posts will do better. Chances are that owing to their social and cultural handicaps they may be generally a shade less competent. But, on the other hand, they will have the great advantage of possessing first hand knowledge of the sufferings and problems of the backward sections of society. This is not a small asset for field workers and policy makers even at the highest level.

13.7 It is no doubt true that the major benefits of reservation and other welfare measures for Other Backward Classes will be cornered by the more advanced sections of the backward communities. But is not this a universal phenomenon ? All reformist remedies have to contend with a slow recovery along the hierarchical gradient; there are no quantum jumps in social reform. Moreover, human nature being what it is, a "new class' ultimately does emerge even in classless societies. The chief merit of reservation is not that it will introduce egalitarianism amongst OBCs when the rest of the Indian society is seized by all sorts of inequalities. But reservation will certainly erode the hold of higher castes on the services and enable OBCs in general to have a sense of participation in running the affairs of their country.

13.8 It is certainly true that reservation for OBCs will cause a lot of heart buring to others. But should

57

Notes: Reproduction from official government archival material for the Backward Classes Commission (Mandal) Report (first part, English). Left panel shows the report cover page; right panel shows the recommendations chapter page discussing OBC reservations and social/educational backwardness. Source: National Commission for Backward Classes archive (https://ncbc.nic.in/Writereaddata/Mandal%20Commission%20Report%20of%20the%201st%20Part%20English635228710201364644.pdf), accessed February 19, 2026.

Why *Indra Sawhney* is a landmark for identification logic. The Supreme Court's judgment in *Indra Sawhney v. Union of India* (1992), often called the "Mandal case," is pivotal because it both upheld backward-class identification for reservations and clarified the conceptual bridge between caste and class in Indian social reality (Supreme Court of India, 1992). It also introduced creamy-layer exclusion for OBC benefits, which is useful for interpretation in this paper: even where reservation eligibility changes, class-relevant social and cultural disadvantage can persist. This is the legal basis for treating OBC creamy-layer outcomes as informative rather than anomalous in hiring analyses.

A.2 Scholarly Foundations for a Class-Relevant Mechanism

Varna and jāti are not interchangeable. Indian scholarship has long emphasized a distinction often blurred in English-language usage. Basham notes that "caste" has been applied

indiscriminately to both *varna* and *jāti*, and calls that usage false terminology (Basham, 1954). In this distinction, *varna* is a broad classificatory order, while *jāti* refers to localized endogamous communities. Béteille similarly treats *jāti* as closer to ethnic grouping in many contexts (Béteille, 1996). For the present paper, this distinction matters because state administration (SC/ST/OBC) aggregates many local groups into policy categories organized around social and educational disadvantage; the empirical object is therefore an administratively aggregated structure, not an attempt to map each local *jāti* into a canonical four-*varna* schema.

Class theory clarifies the interview mechanism. Class theory predicts that evaluators infer social position through bundles of capitals revealed in interaction. In Weberian terms, market position and status closure are analytically distinct but empirically intertwined (Weber, 1978a,b; Parkin, 1979). In Bourdieu's framework, durable advantage reflects joint accumulation of economic, social, and cultural capital, especially embodied forms that become legible through speech, ease, demeanor, and biographical narrative (Bourdieu, 1986, 1984; Lareau, 2003; Lamont and Lareau, 1988). Rivera's work on elite hiring in the U.S. documents a closely related "fit" channel in final-round selection (Rivera, 2012). These frameworks jointly predict exactly the stage timing observed in this paper: penalties should emerge where evaluators can probe biography and interaction, not necessarily where candidates are merely visible.

The closure perspective links these interactional judgments to durable inequality. In Parkin's neo-Weberian account, exclusionary closure controls access to valued positions using criteria that appear neutral but systematically advantage incumbents (Parkin, 1979). In contemporary elite labor markets, "fit" can play that role: it converts diffuse judgments about polish, confidence, and cultural ease into gatekeeping decisions without explicit caste naming. The same logic clarifies OBC-CL interpretation. Creamy-layer status is often misread as if crossing an income threshold eliminates structural disadvantage, but cumulative-capital models imply otherwise: inherited educational trajectories, language socialization, and networks can remain unequal even when current income rises.

Major historical and anthropological accounts likewise caution against reducing caste to a visually obvious phenotype. The colonial state played a substantial role in systematizing and reifying categories through censuses and administrative practice (Dirks, 2001; Bayly, 2008; Cohn, 1987). If categories are historically institutionalized and administratively stabilized, discrimination can operate through institutional and interactional channels even when phenotype is weakly informative.

A standard objection is that class is mutable while caste is fixed. The persistence literature suggests this contrast is often overstated in empirical terms. Long-run surname evidence shows high intergenerational persistence for social rank across societies (Clark, 2014), and Indian scholarship documents durable reproduction through education, occupations, networks, and marriage markets even under formal mobility (Deshpande, 2011b; Munshi, 2019; Jodhka, 2015). For the mechanism studied here, the practical implication is straightforward: evaluators can still apply higher interview thresholds based on cumulative social markers.

A.3 Comparative Implications for This Paper

Relative to many race settings where group assignment is more phenotypic at first exposure (Omi and Winant, 2014), the mechanism here is biographical and interactional. That difference helps explain why group debates and personal interviews can diverge sharply despite both being subjective rounds. Group debates may convey some social signals, but personal interviews create structured opportunities to elicit family background, schooling trajectories, cultural references, and rapport cues that are more actionable for inference. In most class-discrimination settings, researchers rely on noisy proxies (income, parental education, neighborhood); India's SC/ST/OBC architecture offers comparatively high-signal administrative coding of cumulative disadvantage, improving empirical identification of where penalties arise in a multi-stage pipeline.

Institutional boundaries and measurement validity. Indian constitutional law distinguishes state obligations from private-sector discretion. Reservation rules bind public education and public employment, but do not impose equivalent quotas on private multinational hiring in this setting (Galanter, 1984; Jodhka, 2008). This boundary is useful for identification: if caste penalties appear where reservation mandates do not directly govern screening decisions, mechanical compliance is less plausible than discretionary late-stage evaluation. Administrative categories are coarse relative to lived heterogeneity, but they remain informative for group-threshold questions; disaggregated results in Appendix B reinforce this point, including for OBC-CL.

The combined legal and sociological framing also clarifies the empirical sequence in Section 6. In settings where category assignment is immediate at first exposure, discrimination should often appear earlier in screening. In this caste setting, where assignment is frequently biographical and probabilistic, penalties should emerge where interviews permit denser social inference. The legal codification of caste categories does not predetermine that timing; it makes the timing test measurable.

Table A.1 is a compact comparative taxonomy, not an exhaustive theory of stratification. It should be read as an analytical map for discrimination mechanisms. Three dimensions are most diagnostic for this paper: (i) primary identification mechanism (phenotypic versus biographical inference), (ii) degree of state codification, and (iii) whether the system bundles multiple forms of capital. On these dimensions, the interview-stage evidence most closely matches a class-relevant mechanism operating through administratively coded caste categories.

Table A.1: Comparing Caste, Class, and Race Across Nine Dimensions Relevant to Discrimination Research

Dimension	Race	Class	Caste
Primary identification	Phenotypic (skin color, facial features)—visible at a glance (Omi and Winant, 2014)	Inferred from credentials, consumption, accent (in extended conversation), networks (Bourdieu, 1986)	Inferred from biographical markers (surname, hometown, parental occupation, educational pedigree) (Thorat and Attewell, 2007)
Intergenerational persistence	High (but boundaries shift historically) (Ignatiev, 1995)	High (0.7–0.8 correlation across 10–15 generations) (Clark, 2014)	High (administratively codified) (Mandal Commission, 1980)
Formal mutability	Generally immutable (though "passing" documented) (Hobbs, 2014)	Formally mutable (though empirically persistent)	Formally immutable (though Sanskritization documented) (Srinivas, 1966)
Endogamy	Historically enforced; declining	Assortative mating by education/status (Schwartz, 2013)	Strong; marriage markets sort explicitly on caste (Banerjee et al., 2013)
State classification	Census categories; affirmative action	Rarely codified administratively	Explicit constitutional categories (SC/ST/OBC) (Galanter, 1984)
Basis of classification	Phenotype + ancestry	Occupation, income, education, credentials	"Socially and educationally backward" (11 composite indicators) (Mandal Commission, 1980)
Genetic distinctiveness	Population structure exists but continuous	N/A	Historic gene flow across boundaries; not phenotypically distinguishable (Moorjani et al., 2013)
Capitals bundled	Some correlation with economic/social capital	Economic, cultural, social, symbolic (Bourdieu, 1986)	Economic, cultural, social, symbolic (administratively codified) (Deshpande, 2011b)
Discrimination mechanism	Name-based identification → taste-based or statistical (Bertrand and Mullainathan, 2004)	Biographical inference → "fit" screening (Rivera, 2012)	Biographical inference → "fit" screening

Note: This table synthesizes scholarship across sociology, economics, anthropology, and legal studies to compare three stratification systems across nine dimensions relevant to discrimination research. The claim is not that caste *is* class in some essential sense—caste encompasses ritual, religious, and endogamous dimensions that have no direct parallel in Western class systems. Rather, the table illustrates why the class framework is most analytically productive for understanding the *discrimination mechanism* documented in this paper. Three patterns are central: (1) The *identification mechanism* aligns caste with class: both require inference from biographical markers (credentials, accent, hometown, parental occupation), whereas race in the U.S. context is identified phenotypically at a glance. This explains why caste penalties in our data emerge at interviews—where biographical probing occurs—not at group debates where candidates are visually observed. (2) India's *administrative codification* of caste through socioeconomic criteria (the Mandal Commission's 11 indicators) provides a measurement advantage unavailable for class in Western contexts: the state has bundled economic, cultural, and social capitals into SC/ST/OBC categories. (3) All three systems exhibit high *intergenerational persistence*—the assumption that class is "mutable" while caste is "immutable" is not well aligned with persistence evidence (Clark, 2014). Citations in each cell indicate key scholarly sources. See Section 2 for the main text treatment and Appendix A for extended discussion.

B Robustness to Category Aggregation

Motivation, setup, and key takeaway. This section tests whether the pooled disadvantaged-category result is masking composition effects. The core takeaway is that interview-stage penalties remain large across disaggregated categories, including OBC creamy-layer (OBC-CL), which is reservation-ineligible. If reservation eligibility or current income were the main channel, OBC-CL penalties should attenuate sharply. The tables compare each category against the General/Unreserved reference group using the main controls and fixed effects. Across hiring, job quality, and post-hire outcomes, disaggregation confirms that pooling does not fabricate the main result. Interview-stage penalties remain strong for OBC-CL despite reservation ineligibility, which is more consistent with a screening mechanism loading on cumulative social position than on current income alone.

Results: Hiring Success Rates. Table B.1 reports stage-by-stage success-rate gaps. Column (1) reproduces the pooled disadvantaged-category result; Columns (2)–(5) disaggregate to SC, ST, OBC non-creamy-layer, and OBC-CL.

Table B.1: Success Rate Differences by Disaggregated Caste Category

	Coefficient on Category Indicator (Relative to General)				
Stage	Disadv. Caste	SC	ST	OBC–NCL	OBC–CL
App → Test	-0.0025	-0.0254***	-0.0396***	0.0112***	-0.0511***
	(0.0018)	(0.0019)	(0.0030)	(0.0023)	(0.0054)
Test → GD	-0.0002	-0.0189***	-0.0584***	0.0088***	-0.0526***
	(0.0023)	(0.0036)	(0.0056)	(0.0022)	(0.0063)
GD → Int	-0.0005	0.0042	-0.0389***	-0.0002	0.0074
	(0.0051)	(0.0056)	(0.0069)	(0.0053)	(0.0117)
Int → Offer	-0.2938***	-0.2799***	-0.2704***	-0.2958***	-0.3634***
	(0.0106)	(0.0117)	(0.0143)	(0.0107)	(0.0204)
Observations by stage: App→Test: 305,539; Test→GD: 183,143; GD→Int: 148,480; Int→Offer: 38,342					
Year FE	✓	✓	✓	✓	✓
Job FE	✓	✓	✓	✓	✓
Other Controls	✓	✓	✓	✓	✓

Notes: This table reports coefficients from regressions of stage-specific hiring success on caste category indicators, with General category as the omitted reference group. Each row represents a different hiring stage, conditioning on candidates who reached that stage. "App → Test" conditions on all applicants; "Test → GD" conditions on passing the written test; "GD → Int" conditions on passing the group debate; "Int → Offer" conditions on reaching the final interview. Column 1 pools all disadvantaged categories (SC + ST + OBC-NCL + OBC-CL) into a single indicator, replicating the main analysis specification. Columns 2–5 disaggregate by specific category. All specifications include job fixed effects, cohort-year fixed effects, and individual controls (college GPA, entrance exam score percentile, 10th grade marks, 12th grade marks, and prior internship sector). Department and degree stream fixed effects are also included. Standard errors clustered at the job level appear in parentheses. The Observations row reports the sample size at each stage; sample sizes decrease as candidates are eliminated (App→Test: 305,539; Test→GD: 183,143; GD→Int: 148,480; Int→Offer: 38,342). Significance levels: * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Column (1) shows the pooled pattern from the main text: negligible pre-interview gaps and a large interview penalty. Disaggregated estimates show large interview penalties across disadvantaged categories, including OBC-CL, which is difficult to reconcile with current-income-only accounts.

Results: Job Quality. Table B.2 reports job-quality differences (log mean salary in USD PPP) using the same category split. The salary patterns mirror hiring-success results, including sizable OBC-CL penalties.

Table B.2: Job Quality Differences by Disaggregated Caste Category

	Coefficient on Category Indicator (Relative to General)				
Outcome	Disadv. Caste	SC	ST	OBC–NCL	OBC–CL
Log Salary (Offers)	-0.1049***	-0.0954***	-0.1150**	-0.1039***	-0.1582***
	(0.0132)	(0.0266)	(0.0513)	(0.0140)	(0.0587)
Log Salary (Choices)	-0.1035***	-0.0682*	-0.1280*	-0.1075***	-0.1377*
	(0.0174)	(0.0351)	(0.0677)	(0.0185)	(0.0775)
Observations	2,527	2,527	2,527	2,527	2,527
	2,527	2,527	2,527	2,527	2,527
Year FE	✓	✓	✓	✓	✓
Department FE	✓	✓	✓	✓	✓
Other Controls	✓	✓	✓	✓	✓

Notes: This table reports coefficients from regressions of log mean salary (USD PPP) on caste category indicators, with General category as the omitted reference group. "Log Salary (Offers)" uses mean salary across all jobs where the student received an offer. "Log Salary (Choices)" uses the mean salary of jobs the student accepted. Column 1 pools all disadvantaged categories; Columns 2–5 disaggregate. Controls include college GPA, entrance exam score, 10th/12th grade marks, and prior internship sector. Department, stream, and year fixed effects included. Standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Results: Post-Hire Outcomes. Table B.3 extends the same disaggregation to promotions and first-job exit.

Table B.3: Post-Hire Outcomes by Disaggregated Caste Category

	Coefficient on Category Indicator (Relative to General)				
Outcome	Disadv. Caste	SC	ST	OBC–NCL	OBC–CL
Promotion	0.0962***	0.0805	0.0349	0.0967***	0.1760**
	(0.0266)	(0.0544)	(0.0796)	(0.0270)	(0.0731)
Left Job	-0.1204***	-0.1149**	-0.1814*	-0.1111***	-0.2916***
	(0.0300)	(0.0495)	(0.0962)	(0.0332)	(0.0734)
Observations	2,298	2,298	2,298	2,298	2,298
	2,298	2,298	2,298	2,298	2,298
Year FE	✓	✓	✓	✓	✓
Job FE	✓	✓	✓	✓	✓
Other Controls	✓	✓	✓	✓	✓

Notes: This table reports coefficients from regressions of post-hire outcomes on caste category indicators, with General category as the omitted reference group. "Promoted" indicates receiving a promotion within the observation window. "Left First Job" indicates departing the initial employer. Column 1 pools all disadvantaged categories; Columns 2–5 disaggregate. Controls include college GPA, entrance exam score, 10th/12th grade marks, prior internship sector, job fixed effects, department fixed effects, degree stream fixed effects, and cohort-year fixed effects. For promotion regressions, job tenure is also included as a control. Standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

The post-hire results remain consistent with a higher-bar mechanism: disadvantaged hires, including OBC-CL, are more likely to be promoted and less likely to leave conditional on being hired.

C Caste in Corporate DEI Policy

This appendix asks whether caste appears in formal private-sector diversity policy. I combine qualitative evidence on employer attitudes with a systematic content analysis of DEI pages from firms recruiting in this market.

Qualitative evidence on employer attitudes. India-specific scholarship documents sustained employer resistance to caste-conscious private-sector affirmative action, often framed in the language of merit and efficiency (Jodhka, 2008; Jodhka and Newman, 2007; Thorat and Newman, 2010). In an invited review of corporate responses to reservation debates, Jodhka (2008) summarizes the stance as follows:

> *"The response of the Indian private sector to the proposal was unequivocally negative... Reservations, they repeatedly argued, would work against 'merit' and 'efficiency'."*

The same review also emphasizes the gap between formal merit rhetoric and how candidates are screened in practice:

> *"The suitability of a candidate is rarely judged on his or her formal qualification alone... one of the most important questions they ask their candidates during the interviews is about their family background."*

Related evidence from field experiments and labor-market studies is consistent with this tension between formal universalist language and caste-linked treatment in practice (Thorat and Attewell, 2007; Cassan et al., 2021; Munshi, 2019). These studies motivate the quantitative test below: whether caste is visible in the public DEI language of firms active in elite recruitment.

Data and collection protocol. The sample includes 1,017 firms recruiting at elite Indian engineering colleges, spanning 377 India-headquartered and 640 foreign-headquartered firms across technology, consulting, manufacturing, finance, consumer, and pharmaceutical sectors. For each firm, I queried standard DEI paths (for example, `/diversity`, `/dei`, `/inclusion`) and India-specific paths (`.in`, `.co.in`, `/india/`) when available, plus closely related alternatives such as `/belonging` and `/culture`. Pages were parsed to remove navigation/scripts and retained if substantive text exceeded 200 words. This yielded 191 qualifying DEI pages; 161 have clean headquarters metadata (59 India-headquartered, 102 foreign-headquartered), while 30 pages are used in aggregate counts but excluded from HQ splits.

Coding dictionary. Each DEI page was scanned for mentions of demographic categories using a fixed keyword dictionary. The coding rule is intentionally permissive at detection (any keyword hit) and conservative at interpretation: a detected mention is treated as evidence of public salience, not of substantive policy depth or implementation.

Note: A page is coded as mentioning a category if any corresponding keyword appears in the extracted DEI text.

Category	Keywords
Gender	women, gender, female, male, non-binary, gender equality, gender parity
Race/Ethnicity	race, racial, ethnicity, ethnic, minorities, black, asian, hispanic, latino
LGBTQ+	lgbtq, lgbt, lesbian, gay, bisexual, transgender, queer, sexual orientation
Caste	**caste, scheduled caste, scheduled tribe, sc/st, dalit, obc, other backward class**
Socioeconomic	socioeconomic, economic background, class, low-income, first-generation
Disability	disability, disabled, differently abled, accessibility, pwd, handicapped
Age	age, generational, ageism, older workers, age diversity, multigenerational
Religion	religion, religious, faith, beliefs, religious diversity

Historical persistence check. To test persistence rather than a one-year snapshot artifact, I collected Wayback Machine DEI snapshots for 2017–2024 (262 archived pages). Caste mentions remain rare over time: 1.9% of snapshots mention caste, with zero mentions in 2017, 2018, 2023, and 2024.

Measurement notes and denominator discipline. All percentages in Figure C.1 are computed over firms with qualifying DEI text and known headquarters metadata. The 826 firms with no qualifying DEI page are excluded from figure denominators by construction but explicitly reported in the notes because their absence is itself substantively informative for interpretation.

Main finding. Figure C.1 shows that caste is largely absent from formal DEI language. Among firms with substantive DEI pages, caste appears in only 5.1% of India-headquartered pages and 1.0% of foreign-headquartered pages.

Interpretation and limits. Two facts drive interpretation. First, 81% of sampled firms have no qualifying DEI page at all (826 of 1,017 firms). Second, among firms with substantive DEI content, caste is rarely mentioned (5.1% of India-headquartered pages and 1.0% of foreign-headquartered pages). Together, these patterns indicate low caste salience in formal policy language and make formal caste-conscious hiring or promotion initiatives an unlikely explanation for the promotion advantage in the main results.

This interpretation is aligned with the India-specific qualitative literature above, which documents strong resistance to explicit caste-conscious policy in the corporate sector even when screening practices themselves can remain socially selective (Jodhka, 2008; Jodhka and Newman, 2007; Thorat and Newman, 2010). At the same time, the DEI-page evidence is about *public formal policy language*, not private internal practice. It does not prove that firms never use informal channels; it narrows one alternative explanation by showing that explicit caste-targeted DEI commitments are uncommon in the firms and period studied.

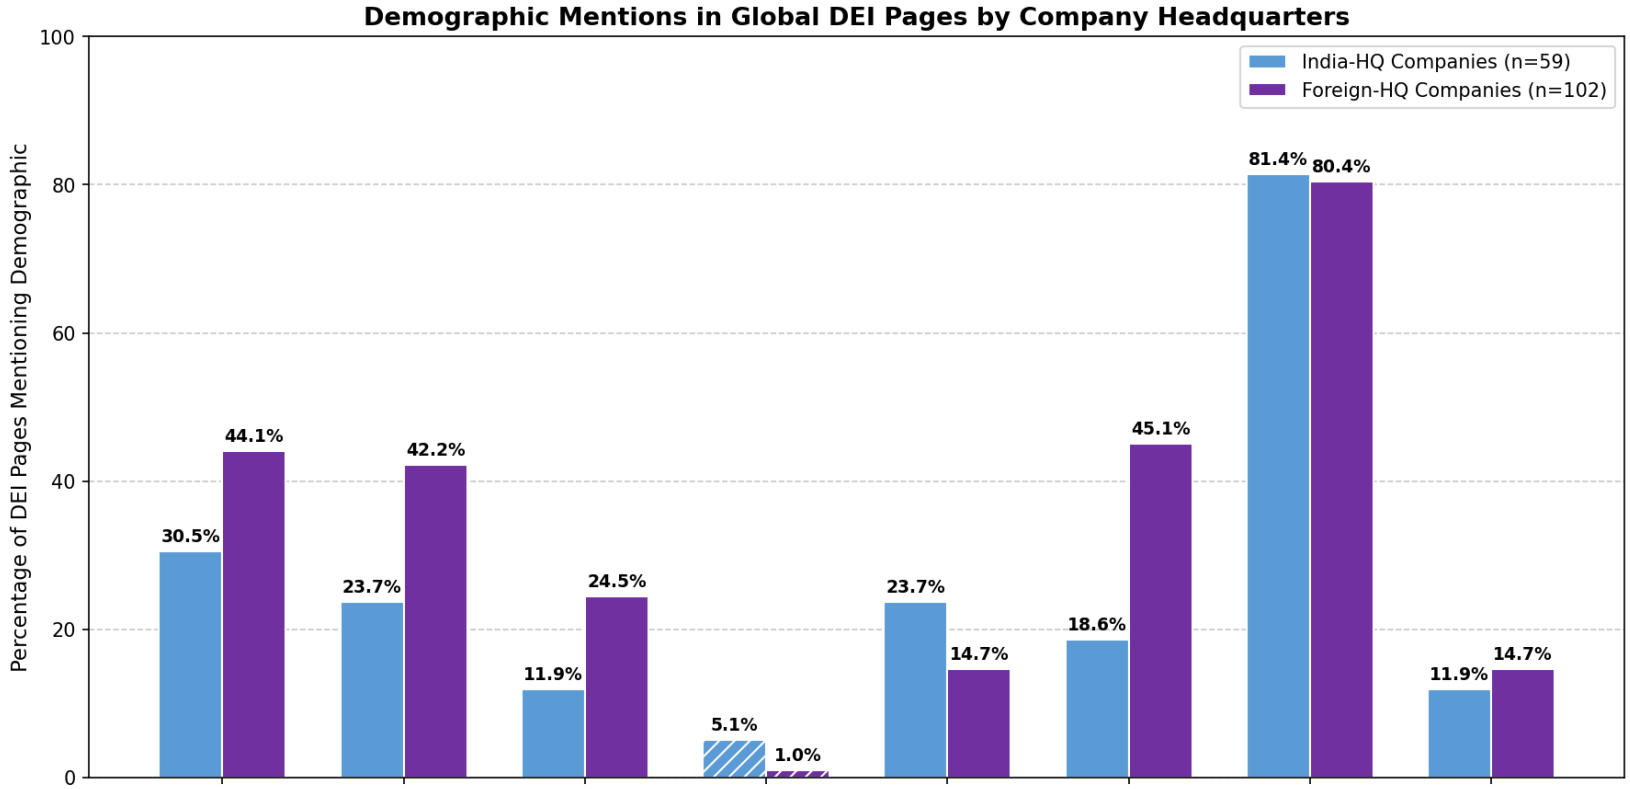


Figure C.1: **Demographic Mentions in Corporate DEI Statements by Company Headquarters**

Notes: Figure C.1 shows the percentage of corporate DEI pages that mention each demographic category, separately by headquarters location. We scraped 1,017 multinational companies recruiting at elite Indian engineering colleges; 191 had DEI pages with substantive content (>200 words), of which 161 could be matched to headquarters (59 India-headquartered, 102 foreign-headquartered). Each bar is the share of DEI pages mentioning the category. Caste keywords searched: caste, scheduled caste, scheduled tribe, SC/ST, dalit, OBC, other backward class. The 826 companies (81%) with no qualifying DEI page are not shown.

D Representativeness of the Job Fair at the Elite College

This appendix demonstrates that the studied on-campus job fair is representative of elite hiring in India. We present two forms of evidence: (1) systematic quantitative analysis of placement processes at 108 elite Indian institutions, and (2) corroborating observations from prior research and institutional documentation.

D.1 Systematic Evidence of Standardization Across Elite Institutions

To quantify the degree of standardization in campus placement, we systematically examined placement processes at 108 elite Indian institutions—including all 23 Indian Institutes of Technology (IITs), all 31 National Institutes of Technology (NITs), 20 Indian Institutes of Management (IIMs), and 34 other elite institutions (IIITs, private engineering colleges, and university colleges)—using automated web scraping of official placement office websites in November 2025.

Of these 108 institutions, 28 (25.9%) published detailed placement procedures. Transparency rates varied by institution type: IIMs showed highest transparency (50%), followed by IITs (26.1%) and NITs (16.1%). The 74% of institutions that did not publish detailed rules likely follow similar practices, as those that do publish reveal nearly identical processes.

Among institutions with detailed documentation, placement processes showed remarkable uniformity. Where institutions described their screening procedures, they consistently reported the same four-stage pipeline documented in this paper: application screening, written aptitude tests, group debates, and personal interviews. This uniformity extends beyond the IIT system to include NITs, IIMs, and private engineering colleges, confirming that the recruitment pipeline reflects a nationwide institutional norm rather than an idiosyncrasy of the studied college.

Interpretation and limit. The 28 institutions with detailed procedures are not a random sample of all colleges; they are the subset with transparent placement documentation. The representativeness argument therefore combines two pieces of evidence: near-identical procedures among transparent institutions, and strong institutional incentives for conformity in campus recruitment logistics (shared recruiters, shared timelines, shared placement-office templates). This does not imply perfect institutional homogeneity, but it supports using the focal college as a high-information case of the broader elite placement architecture.

D.2 Corroborating Evidence

Four additional observations, drawn from prior research and institutional documentation, corroborate the systematic evidence above:

1. **Historical origins of standardization.** 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. This historical pattern explains the convergence documented in our systematic analysis.

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, the systematic quantitative evidence and corroborating observations demonstrate that the focal job fair closely mirrors elite hiring pathways across India's top universities. The four-stage screening pipeline documented in this paper is not unique to the studied institution but represents standard practice across India's elite higher education landscape.

E Employment-Platform Coverage Benchmarks and Profile-Matching Methodology

This appendix benchmarks the coverage of our professional-profile matching. External evidence suggests LinkedIn reaches roughly 80% of U.S. bachelor's-degree holders, with higher penetration in STEM (Wilmers et al., 2025); LinkedIn's own STEM report documents very high platform engagement among STEM master's members (Baird et al., 2023). Cross-country evidence shows similar overrepresentation of STEM professionals on the platform, including in India (Kashyap and Verkroost, 2021). Industry benchmarks are consistent: IT workers are about 2.5 times overrepresented relative to the U.S. labor force (Tambe et al., 2021), and elite-university alumni audits report coverage above 95% after de-duplication (Zhang and Vucetic, 2016). In our data, matches are validated against placement records (firm, year, title), manually audited for ambiguous names, and de-duplicated for likely collisions. Taken together, these checks make a 98% match rate plausible for an elite STEM cohort.

E.1 Career Outcome Data Collection and Professional Profile Matching Methodology

Post-hire outcomes (promotions and tenure) are measured by matching alumni records to professional profiles on a major employment-network platform. Coverage is high in comparable populations: platform studies estimate roughly 80% coverage for U.S. bachelor's-degree holders, with substantially higher penetration in STEM and elite-college cohorts (Tambe and Hitt, 2012; Wilmers et al., 2025). Our setting combines that baseline with unusually rich institutional data, producing a near-census match rate.

Matching begins with curated alumni lists (names, cohorts, majors) for direct searches cross-checked to placement records. For missing cases, institutional roll-number patterns narrow candidate sets, then matches are verified using job history, graduation year, and placement data. Operationally, the process runs in three passes: (i) deterministic matches on exact name+cohort+firm anchors; (ii) probabilistic expansion for common names filtered by degree-year-major compatibility; and (iii) manual adjudication of ambiguous cases using title progression and employer history. Manual audits and de-duplication (including name-variant collapse) further reduce false matches. With placement-office records providing quasi-unique identifiers, the protocol delivers a 98% match rate with high confidence in identity resolution.

F HR Professional Interviews: Theme-Ordered Excerpts

To complement the quantitative evidence, I conducted semi-structured interviews with five senior HR professionals involved in campus recruitment at major multinationals: Tata Consultancy Services, Tata Motors, Aditya Birla Group, IRIS Software Group, and Jindal South West Steel.[18] The interviewees collectively represent over 70 years of campus-recruitment experience across IT services, enterprise software, manufacturing, steel/metals, and diversified conglomerates. This is a small convenience sample, so I use it for institutional triangulation rather than stand-alone identification.

Interview context. Interviews were conducted in 2024 and focused on: (i) pool hiring logistics, (ii) who conducts personal interviews, (iii) what personal interviews probe, (iv) whether interviewers later supervise hired candidates, and (v) how talent-acquisition teams are evaluated.

Theme 1: Retention and promotion are treated as quality metrics. Interviewees described first-job retention and internal growth as key indicators of whether campus hiring was effective, and described promotion as performance-linked rather than automatic.

> *"It's evaluated on the quality of the overall campus hires and their retention and their cultural fit and their growth in the organization."* (HR head, Aditya Birla Group)

> *"Performance, it is basically the retention part... we invest a lot of time and money on them... if we lose them, then the entire thing is lost."* (HR head, Tata Motors)

> *"[A]fter six months... we do a thorough evaluation and then we decide who are those top performers... [f]or promotion... some will get, some will not get."* (HR head, IRIS Software Group)

> *"[Past] performance rating... feeds into your promotion... [t]he first promotion kicks in minimum three years after."* (HR head, Tata Consultancy Services)

> *"The more important factor that we look at is the attrition from the pool that we have hired... we do look at second year, third year performance as a data point, along with attrition."* (HR head, Jindal South West Steel)

Theme 2: Central pool hiring creates separation from eventual supervisors. Interviewees consistently described campus recruitment as hiring into a central pool, with eventual team/location assignment happening after offers and with weak direct alignment between interviewers and final supervisors.

> *"Predominantly campus hiring will be for a central pool. It is not for any specific business."* (HR head, IRIS Software Group)

[18] Names are anonymized for confidentiality; firm names are retained for institutional context.

"So it's a pool which is hired... which person will go to which location is not very firmed up... it's not the direct manager who is interviewing them." (HR head, Tata Motors)

"[Graduate engineer trainees] are hired by a pool of panelists... placement in a specific function is mix and match after offers have been made." (HR head, Jindal South West Steel)

"Those line managers [who interview] are the ones who are available at that time... There is no line manager to student alignment... So it's almost like a lottery." (HR head, Aditya Birla Group)

"The placements are so random, that it's almost like lottery. Accidentally, you can get one person you may have interviewed, but there's no interviewer matching done when placement happens." (HR head, Tata Consultancy Services)

Theme 3: Process metrics dominate, with weak downstream feedback. Interviewees emphasized recruiter evaluation on turnaround time and conversion metrics, with limited linkage to later appraisal and promotion outcomes. Some firms track aggregate cohort outcomes (for example retention), but this is distinct from systematic interviewer-level accountability for specific hire performance.

"For talent acquisition team, it's mainly the TAT [turnaround time]... how much time they took to hire... and how many accepted offers really joined. Their performance is not linked to the performance of the hired candidates." (HR head, Tata Motors)

"It is almost impossible to track back their performance, their appraisal or their promotion to the recruitment branch." (HR head, Tata Consultancy Services)

"After six months of probation we do an evaluation... I do not think too much of data goes to recruitment after six months... beyond that, the ownership will be with the business." (HR head, IRIS Software Group)

"It's evaluated on the quality of the overall campus hires and their retention and their cultural fit and their growth in the organization." (HR head, Aditya Birla Group)

"There is no qualitative parameters to assess the hiring process itself, except looking at the offer drop ratio." (HR head, Jindal South West Steel)

Theme 4: Biographical probing is embedded in fit assessment. Interviewees described personal rounds as open-ended conversations where family background, schooling history, communication style, hobbies, and broader background frequently enter screening.

"In the HR round, we do ask for family backgrounds... we ask about their education backgrounds, where they have studied... and what kind of background they come from." (HR head, Tata Motors)

> *"[When discussing communication expectations, some candidates] came from backgrounds where their parents were working as domestic help or driver... you can't expect them to read encyclopedias and Enid Blyton."* (HR head, Tata Consultancy Services)

> *"[Do things like hobbies and schooling come up in conversation?] Yes, it does. It does for sure."* (HR head, IRIS Software Group)

> *"These are more aptitude and attitude and aspiration related questions... what's the fit in the organization, that's the basic... [and when asked if interviewers can ask about family background and schooling:] there are no issues on any questions."* (HR head, Aditya Birla Group)

> *"[In personal rounds, interviewers ask about] hobbies, extracurricular[s], [and] what [students] did beyond [the] curriculum... those are more of building rapport [and] trying to understand and gauge... [and] those basically help us take a call."* (HR head, Jindal South West Steel)

Taken together, these interviews are consistent with the transcript evidence in Section 6.4: personal interviews are a setting where biographical markers become salient, and interviewer-supervisor separation plus weak performance feedback can allow biased fit screening to persist.

Comparative literature anchor. These theme-ordered excerpts also align with evidence from elite hiring outside India, where final-round assessments often weight rapport and cultural matching heavily relative to formal credentials (Rivera, 2012, 2015; Laurison and Friedman, 2019). The external literature does not identify caste-specific channels; it helps contextualize why biographical probing and fit judgments can become decisive in discretionary final rounds.

G Robustness to Functional Form

A potential concern is that our main specifications impose linearity on ability controls (GPA, entrance exam scores, 10th and 12th grade marks). If the true relationship between these controls and outcomes is nonlinear, the caste coefficient could be biased. This section addresses this concern by comparing our main (linear) specification against three flexible alternatives: natural splines, decile indicators for each continuous ability measure, and saturated department $\times$ year $\times$ stream fixed effects. We examine robustness for success rates (Section G.1), job quality (Section G.2), post-hire outcomes (Section G.3), and surname-split analyses (Sections G.4 and G.5).

G.1 Success Rates

Table G.1 reports the coefficient β on disadvantaged caste from Equation 1, estimated at each hiring stage under different functional form specifications. The natural spline specification uses four degrees of freedom for each continuous ability control—a standard choice that provides sufficient flexibility to capture nonlinearities while avoiding overfitting (Harrell, 2015). The decile specification replaces each continuous control (GPA, entrance exam score, 10th grade marks, 12th grade marks) with ten indicator variables corresponding to deciles of the variable's distribution, allowing arbitrary nonlinear relationships.

The results demonstrate remarkable stability. At the critical interview-to-offer stage, the caste coefficient ranges from -0.293 (decile bins) to -0.298 (saturated FE), compared to -0.294 in the main specification—a difference of less than 1.5%. The early-stage coefficients remain small and economically negligible across all specifications, though some reach marginal statistical significance under flexible specifications. This stability indicates that our findings are not artifacts of functional form assumptions.

Table G.1: **Success Rate Robustness: Functional Form Specifications**

	App → Test (1)	Test → Debate (2)	Debate → Int (3)	Int → Offer (4)
Panel A: Main Specification (Linear Controls)				
Disadv. Caste	-0.003*	-0.000	-0.000	-0.294***
	(0.001)	(0.002)	(0.003)	(0.005)
N	305,539	183,143	148,480	38,342
Panel B: Natural Splines (df=4)				
Disadv. Caste	0.018***	0.015***	-0.003	-0.295***
	(0.001)	(0.002)	(0.003)	(0.005)
N	305,539	183,143	148,480	38,342
Panel C: Decile Bins				
Disadv. Caste	0.012***	0.011***	-0.001	-0.293***
	(0.001)	(0.002)	(0.003)	(0.005)
N	305,539	183,143	148,480	38,342
Panel D: Saturated FE (Dept × Year × Stream)				
Disadv. Caste	-0.002	-0.001	-0.001	-0.298***
	(0.001)	(0.002)	(0.003)	(0.005)
N	305,539	183,143	148,480	38,341

Notes: Table G.1 reports the coefficient on disadvantaged caste (SC/ST/OBC indicator) from linear probability models of advancing to the next hiring stage, estimated separately for each stage. The dependent variable is an indicator for advancing. Panel A uses the main specification with linear controls: college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, prior internship sector, department fixed effects, degree stream fixed effects, and cohort-year fixed effects. Panel B replaces the four ability measures with natural splines (df=4). Panel C replaces them with decile indicators. Panel D omits the four ability measures entirely and replaces the separate department, stream, and year fixed effects with saturated department × year × stream fixed effects (i.e., a full set of triple-interaction dummies), so that all ability variation is absorbed nonparametrically within cells. All specifications include job fixed effects. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

G.2 Job Quality

Table G.2 applies the same functional form robustness checks to the job quality analysis. The dependent variable is log average salary of jobs in contention at the offer stage, estimated using Equation 2. Because job quality regressions do not include job fixed effects (salaries are standardized within jobs), we report results for the main specification, natural splines, and decile bins.

The results again demonstrate stability. The caste coefficient on job quality at the offer stage ranges from -0.089 (natural splines) to -0.105 (main specification), with decile bins yielding -0.103. These patterns confirm that functional form assumptions do not drive the job quality findings.

Table G.2: **Job Quality Robustness: Functional Form Specifications**

	App (1)	Test (2)	Debate (3)	Interview (4)	Offer (5)	Choice (6)
Panel A: Main Specification (Linear Controls)						
Disadv. Caste	0.004	−0.009	−0.006	−0.018*	−0.105***	−0.103***
	(0.005)	(0.007)	(0.007)	(0.010)	(0.013)	(0.017)
N	4,005	3,618	3,266	3,031	2,527	2,527
Panel B: Natural Splines (df=4)						
Disadv. Caste	0.005	0.000	0.004	−0.008	−0.089***	−0.085***
	(0.005)	(0.007)	(0.007)	(0.010)	(0.013)	(0.017)
N	4,005	3,618	3,266	3,031	2,527	2,527
Panel C: Decile Bins						
Disadv. Caste	0.003	−0.003	0.001	−0.017*	−0.103***	−0.103***
	(0.005)	(0.007)	(0.007)	(0.010)	(0.013)	(0.017)
N	4,005	3,618	3,266	3,031	2,527	2,527

Notes: Table G.2 reports the coefficient on disadvantaged caste (SC/ST/OBC indicator) from regressions of log average salary at the offer stage on caste and controls, following Equation 2. Panel A uses the main specification with linear controls: college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, prior internship sector, department fixed effects, degree stream fixed effects, and cohort-year fixed effects. Panel B replaces the four ability measures with natural splines (df=4). Panel C replaces them with decile indicators. Job fixed effects are omitted because the dependent variable is at the student level (an average across jobs), making job identifiers undefined. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

G.3 Post-Hire Outcomes

Table G.3 applies the same functional form robustness checks to post-hire outcomes. Panel A reports promotion (binary indicator equal to 1 if the individual received a promotion), and Panel B reports leaving first job (binary indicator equal to 1 if the individual left their first employer). Both outcomes are estimated with job and year fixed effects plus controls.

The results demonstrate stability. The caste coefficient on promotion ranges from 0.096 (main specification) to 0.117 (decile bins), and the coefficient on leaving first job ranges from −0.170 (main specification) to −0.177 (splines and decile bins). These patterns confirm that the post-hire advantage for disadvantaged caste graduates is not an artifact of functional form assumptions.

Table G.3: **Post-Hire Robustness: Functional Form Specifications**

	Promoted (1)	Left Job (2)
Panel A: Main Specification (Linear Controls)		
Disadv. Caste	0.096***	-0.170***
	(0.026)	(0.027)
N	2,489	2,489
Panel B: Natural Splines (df=4)		
Disadv. Caste	0.104***	-0.177***
	(0.027)	(0.028)
N	2,489	2,489
Panel C: Decile Bins		
Disadv. Caste	0.117***	-0.177***
	(0.027)	(0.028)
N	2,489	2,489

Notes: Table G.3 reports the coefficient on disadvantaged caste (SC/ST/OBC indicator) from post-hire outcome regressions. Column (1), Promoted (Panel A), and Column (2), Left Job (Panel B), are estimated on the same sample. Row groups A–C apply the main linear-controls specification, natural splines (df=4), and decile-bin controls, respectively. All specifications include job fixed effects, department fixed effects, degree stream fixed effects, and cohort-year fixed effects. Standard errors are in parentheses. *p<0.1, **p<0.05, ***p<0.01.

G.4 Success Rates by Surname Identifiability

Appendix G, Table G.4 applies the same functional form robustness checks to success-rate regressions, reported separately for students with caste-identifiable surnames (Panel A) and non-identifiable surnames (Panel B). This table corresponds to the surname-identifiability analysis in Section 6.4.2.

In Panel A (identifiable surnames), the interview-to-offer penalty ranges from -0.224 (saturated FE) to -0.246 (decile bins), compared to -0.233 in the main specification. Notably, early-stage penalties are substantial and stable: the application-to-test coefficient ranges from -0.166 to -0.184 across specifications, consistent with employers inferring caste from surnames at the application stage.

In Panel B (non-identifiable surnames), the interview-to-offer penalty ranges from -0.297 to -0.302 across specifications, while early-stage coefficients remain near zero or slightly positive. This confirms that the differential timing pattern—penalties early for identifiable surnames, only at interviews for non-identifiable surnames—is robust to functional form assumptions.

Table G.4: **Success Rate Robustness by Surname Identifiability**

Panel A: Identifiable Surnames

	App→Test (1)	Test→GD (2)	GD→Int (3)	Int→Offer (4)
Panel A: Main Specification (Linear Controls)				
Disadv. Caste	-0.184***	0.110***	-0.063***	-0.233***
	(0.004)	(0.005)	(0.007)	(0.021)
N	46,129	21,349	18,839	3,750
Panel B: Natural Splines (df=4)				
Disadv. Caste	-0.166***	0.108***	-0.065***	-0.240***
	(0.003)	(0.005)	(0.007)	(0.022)
N	46,129	21,349	18,839	3,750
Panel C: Decile Bins				
Disadv. Caste	-0.179***	0.096***	-0.061***	-0.246***
	(0.003)	(0.005)	(0.008)	(0.023)
N	46,129	21,349	18,839	3,750
Panel D: Saturated FE (Dept × Year × Stream)				
Disadv. Caste	-0.177***	0.122***	-0.063***	-0.224***
	(0.004)	(0.005)	(0.008)	(0.027)
N	46,129	21,347	18,837	3,745

Panel B: Non-Identifiable Surnames

	App→Test (1)	Test→GD (2)	GD→Int (3)	Int→Offer (4)
Panel A: Main Specification (Linear Controls)				
Disadv. Caste	0.026***	-0.011***	0.008***	-0.298***
	(0.001)	(0.002)	(0.003)	(0.005)
N	259,375	161,737	129,580	34,409
Panel B: Natural Splines (df=4)				
Disadv. Caste	0.048***	0.006***	0.005*	-0.300***
	(0.001)	(0.002)	(0.003)	(0.005)
N	259,375	161,737	129,580	34,409
Panel C: Decile Bins				
Disadv. Caste	0.041***	0.002	0.006**	-0.297***
	(0.001)	(0.002)	(0.003)	(0.005)
N	259,375	161,737	129,580	34,409
Panel D: Saturated FE (Dept × Year × Stream)				
Disadv. Caste	0.028***	-0.012***	0.005*	-0.302***
	(0.001)	(0.002)	(0.003)	(0.005)
N	259,375	161,737	129,580	34,408

Notes: Table G.4 reports the coefficient on disadvantaged caste from linear probability models of advancing to the next hiring stage, estimated separately for each stage. Panel A restricts the sample to students with caste-identifiable surnames (n=619), while Panel B restricts the sample to students with non-identifiable surnames (n=3,386). Within each panel, row group A uses the main specification with linear controls, and row groups B–D apply alternative functional forms (natural splines, decile bins, saturated fixed effects). All specifications include job fixed effects. Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

G.5 Job Quality by Surname Identifiability

Appendix G, Table G.5 applies the same functional form robustness checks to the job-quality analysis separately for students with caste-identifiable surnames (Panel A) and non-identifiable surnames (Panel B). This table corresponds to the surname-identifiability analysis in Section 6.4.2.

In Panel A (identifiable surnames), the job-quality penalty at the offer stage ranges from -0.143 to -0.146 across specifications. In Panel B (non-identifiable surnames), the penalty ranges from -0.082 to -0.100. Both samples show stable results across functional form specifications, confirming that the job-quality findings are robust.

Table G.5: **Job Quality Robustness by Surname Identifiability**

Panel A: Identifiable Surnames

	Offer (1)	Choice (2)
Panel A: Main Specification (Linear Controls)		
Disadv. Caste	-0.146***	-0.107**
	(0.033)	(0.049)
N	378	378
Panel B: Natural Splines (df=4)		
Disadv. Caste	-0.145***	-0.103**
	(0.033)	(0.050)
N	378	378
Panel C: Decile Bins		
Disadv. Caste	-0.143***	-0.098*
	(0.035)	(0.051)
N	378	378

Panel B: Non-Identifiable Surnames

	Offer (1)	Choice (2)
Panel A: Main Specification (Linear Controls)		
Disadv. Caste	-0.100***	-0.106***
	(0.014)	(0.019)
N	2,149	2,149
Panel B: Natural Splines (df=4)		
Disadv. Caste	-0.082***	-0.085***
	(0.014)	(0.018)
N	2,149	2,149
Panel C: Decile Bins		
Disadv. Caste	-0.092***	-0.099***
	(0.014)	(0.019)
N	2,149	2,149

Notes: Table G.5 reports the coefficient on disadvantaged caste from regressions of log average salary on caste and controls. Panel A restricts the sample to students with caste-identifiable surnames (n=378 job-getters); Panel B restricts to students with non-identifiable surnames (n=2,149 job-getters). Column (1) reports offer-stage salary and Column (2) reports final job-choice salary. Within each panel, row group A uses the main specification with linear controls, and row groups B–C apply alternative functional forms (natural splines and decile bins). Standard errors are in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

H Robustness to Omitted Variable Bias: Oster Bounds

While the preceding section addresses whether *functional form* of included controls affects our estimates, this section addresses a distinct concern: whether *omitted variables*—factors we do not observe—could explain the caste penalty. We implement the Oster (2019) bounds methodology, which asks: how much would selection on unobservables need to exceed selection on observables to explain away the estimated effect?

The method requires two parameters. First, δ measures the relative degree of selection on unobservables versus observables: $\delta = 1$ means unobservables are equally predictive of outcomes as observables, while $\delta > 1$ means unobservables are *more* predictive. Second, $R^*_{\max}$ is the hypothetical R^2 from a regression including all possible controls. Following Oster (2019), we set $\delta = 1$ and $R^*_{\max} = 1.3\tilde{R}$ (where $\tilde{R}$ is the R^2 from our controlled regression). These choices are conservative: we already control for qualifications directly used in hiring (GPA, entrance scores, experience), making it implausible that unobservables would be more predictive. As Stansbury and Rodriguez (2024) notes in similar discrimination contexts, observable productivity measures are often the very metrics used in selection, limiting scope for unobservables to explain group disparities.

H.1 Success Rates

Table H.1 compares OLS coefficients to Oster-adjusted estimates for success rates at each hiring stage. Results are stable across stages. The interview success coefficient moves only from -0.294 to -0.296. In some specifications (notably test success), the adjusted coefficient moves upward under the Oster calibration. This directional movement reflects the joint behavior of coefficients and R^2 across the short and full specifications under $\delta = 1$; by itself, it does not identify the true sign of unobservables. For stages with near-zero OLS coefficients, δ^* is numerically unstable and is therefore not interpreted.

Table H.1: **OLS vs. Oster-Adjusted Estimates: Success Rates**

	Disadvantaged Caste		
	OLS (1)	Oster-Adjusted (2)	δ^* (3)
App → Test	−0.003*	0.009	0.21
	(0.001)	(0.001)	
N	305,541	305,541	
Test → Debate	−0.001	0.010***	0.05
	(0.002)	(0.002)	
N	183,150	183,150	
Debate → Int	−0.001	−0.003	—†
	(0.003)	(0.003)	
N	148,489	148,489	
Int → Offer	−0.294***	−0.296***	−131.9
	(0.005)	(0.005)	
N	38,378	38,378	

Notes: Table H.1 compares OLS estimates (Column 1) to bias-adjusted estimates (Column 2) using the Oster (2019) methodology. Column (3) reports δ^*, the value of δ that would set the coefficient to zero; negative values indicate selection on unobservables would need to operate in the *opposite direction* of selection on observables. Stage-specific observations are reported in the "N" row beneath each stage. Each row reports the coefficient on disadvantaged caste (SC/ST/OBC indicator) for the probability of advancing from one hiring stage to the next. For each stage, Oster's short specification is the fixed-effects baseline (job and cohort-year FE only), and the full specification adds observed controls. The Oster adjustment uses $\delta = 1$ (selection on unobservables equals selection on observables) and $R^*_{\max} = 1.3\tilde{R}$ (hypothetical R-squared 30% higher than observed). For the interview stage, $\delta^* = -131.9$ indicates the 30 percentage point penalty is highly robust to omitted variable bias. All specifications include job fixed effects, cohort-year fixed effects, and linear controls: college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, prior internship sector indicators, department fixed effects, and degree stream fixed effects. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. †For stages with near-zero OLS coefficients, δ^* is numerically unstable and not reported.

H.2 Job Quality

Table H.2 applies the same Oster bounds methodology to the job quality analysis. The dependent variable is log average salary of jobs in contention at each hiring stage. The caste coefficient at the offer stage moves from −0.105 (OLS) to −0.079 (Oster-adjusted), remaining economically and statistically significant. Under the stated Oster calibration, this stability is difficult to reconcile with omitted-variable bias explaining away the job-quality penalty that emerges at the interview and offer stages.

Table H.2: **OLS vs. Oster-Adjusted Estimates: Job Quality**

	Disadvantaged Caste		
	OLS (1)	Oster-Adjusted (2)	δ^* (3)
Application Stage	0.004	0.005	−1.9
	(0.005)	(0.005)	
N	4,005	4,005	
Written Test Stage	−0.009	0.010	0.47
	(0.007)	(0.008)	
N	3,618	3,618	
Group Debate Stage	−0.006	0.011	0.37
	(0.007)	(0.008)	
N	3,266	3,266	
Interview Stage	−0.018*	0.012	0.62
	(0.010)	(0.011)	
N	3,031	3,031	
Offer Stage	−0.105***	−0.079***	4.0
	(0.013)	(0.014)	
N	2,527	2,527	
Choice Stage	−0.103***	−0.076***	3.7
	(0.017)	(0.019)	
N	2,527	2,527	

Notes: Table H.2 compares OLS estimates (Column 1) to bias-adjusted estimates (Column 2) using the Oster (2019) methodology. Column (3) reports δ^*, the value of δ that would set the coefficient to zero. Stage-specific observations are reported in the "N" row beneath each stage. Each row reports the coefficient on disadvantaged caste (SC/ST/OBC indicator) for log average salary of jobs in contention at each hiring stage. For each stage, Oster's short specification includes cohort-year fixed effects only, and the full specification adds observed controls. The Oster adjustment uses $\delta = 1$ (selection on unobservables equals selection on observables) and $R^*_{\max} = 1.3\tilde{R}$ (hypothetical R-squared 30% higher than observed). For the offer stage, $\delta^* = 4.0$: unobservables would need to be 4 times as predictive as observables to explain away the 10.5% job quality penalty. All specifications include cohort-year fixed effects and linear controls: college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, prior internship sector indicators, department fixed effects, and degree stream fixed effects. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

H.3 Promotion and Tenure Rates

Appendix H, Table H.3 reports Oster bounds for post-hire outcomes in a combined format. Panel A applies the methodology to promotion rates across job selectivity tiers. The promotion advantage for disadvantaged caste graduates is robust to bias adjustment: the coefficient moves from 0.096 (OLS) to 0.093 (Oster-adjusted) in the overall sample, and the positive promotion pattern across selectivity tiers remains.

Panel B applies the same methodology to first-job leaving (tenure). Disadvantaged caste graduates are significantly *less* likely to leave their first job, and this pattern is robust to Oster adjustment. The coefficient moves from −0.170 (OLS) to −0.180 (Oster-adjusted) in the full specification, with $\delta^* = -17.9$: under this calibration, eliminating the estimate would require unobservables to be much stronger than observables and to load in the opposite direction. The tenure results complement the promotion results: disadvantaged hires both show higher promotion and lower first-job leaving, which is hard to reconcile with a pure accurately screened productivity-deficit interpretation.

Table H.3: **OLS vs. Oster-Adjusted Estimates: Post-Hire Outcomes**

Panel A: Promotion Rates

		Disadvantaged Caste		
	N	OLS (1)	Oster-Adjusted (2)	δ^* (3)
Promoted (All)	2,489	0.137***	0.142***	−27.4
		(0.025)	(0.026)	
Promoted (w/ controls)	2,489	0.096***	0.093***	29.6
		(0.026)	(0.028)	
Promoted at Top 75%	2,012	0.104***	0.108***	−27.4
		(0.029)	(0.031)	
Promoted at Top 50%	1,246	0.112***	0.122***	−11.6
		(0.038)	(0.037)	
Promoted at Top 40%	997	0.136***	0.141***	−29.0
		(0.043)	(0.041)	
Promoted at Top 30%	748	0.147***	0.157**	−15.0
		(0.054)	(0.058)	
Promoted at Top 20%	502	0.190**	0.196**	−35.2
		(0.073)	(0.082)	
Promoted at Top 15%	374	0.204**	0.206**	−102.2
		(0.084)	(0.082)	
Promoted at Top 10%	252	0.084	0.074	8.6
		(0.109)	(0.131)	
Promoted (Client Facing)	835	0.107**	0.106***	−1190.9
		(0.044)	(0.039)	
Promoted (Non-Client Facing)	1,654	0.082**	0.075**	11.8
		(0.032)	(0.036)	

Panel B: Job Leaving Rates (Tenure)

		Disadvantaged Caste		
	N	OLS (1)	Oster-Adjusted (2)	δ^* (3)
Overall Specifications				
Left Job (Job FE)	2,489	−0.188***	−0.198	−17.4
Left Job (Job + Year FE)	2,489	−0.186***	−0.197	−17.7
Left Job (Full Controls)	2,489	−0.170***	−0.180	−17.9
By Job Selectivity				
Top 75% Selectivity	2,012	−0.124***	−0.139	−7.9
Top 50% Selectivity	1,246	−0.135***	−0.147	−10.7
Top 40% Selectivity	997	−0.210***	−0.213	−91.9
Top 30% Selectivity	748	−0.189***	−0.200	−17.6
Top 25% Selectivity	623	−0.158**	−0.157	169.6
Top 20% Selectivity	502	−0.152*	−0.128	6.4
Top 15% Selectivity	374	−0.191**	−0.170	9.0
Top 10% Selectivity	252	−0.181	−0.161	8.9
By Client-Facing Status				
Consumer Facing	835	−0.067	−0.070	−21.5
Non-Consumer Facing	1,654	−0.171***	−0.188	−9.7

Notes: Table H.3 compares OLS estimates (Column 1) to Oster-adjusted estimates (Column 2), with δ^* in Column 3. Panel A reports promotion outcomes by selectivity and client-facing status; Panel B reports first-job leaving outcomes by selectivity and client-facing status. For each row, Oster's short specification uses fixed effects only, and the full specification adds observed controls. The Oster adjustment uses $\delta = 1$ and $R^*_{\max} = 1.3\tilde{R}$. Negative δ^* values indicate that unobservables would need to operate in the opposite direction of observables to eliminate the estimated effect. Across both panels, coefficients remain stable after adjustment, supporting robustness of post-hire results to omitted-variable bias. Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

H.4 Job Quality by Surname Identifiability

Appendix H, Table H.4 applies the same methodology separately to students with caste-identifiable surnames (Panel A, 15% of the sample) and non-identifiable surnames (Panel B, 85%). For identifiable surnames, penalties emerge early (written test stage) and remain significant after Oster adjustment. For non-identifiable surnames, pre-interview coefficients are essentially zero, with some Oster-adjusted values moving slightly positive under the calibration assumptions; the penalty still emerges at the offer stage. This pattern reinforces the main text finding that penalties track caste visibility.

Table H.4: **OLS vs. Oster-Adjusted Estimates: Job Quality by Surname Identifiability**

Panel A: Identifiable Surnames

	Disadvantaged Caste		
	OLS (1)	Oster-Adjusted (2)	δ^* (3)
Application	0.012	0.013	−7.9
	(0.014)	(0.016)	
N	619	619	
Written Test	−0.068***	−0.049***	3.3
	(0.019)	(0.021)	
N	552	552	
Group Debate	−0.034*	−0.016	1.8
	(0.018)	(0.022)	
N	500	500	
Interview	−0.105***	−0.078**	3.9
	(0.027)	(0.032)	
N	449	449	
Offer	−0.146***	−0.119***	5.6
	(0.033)	(0.042)	
N	378	378	

Panel B: Non-Identifiable Surnames

	Disadvantaged Caste		
	OLS (1)	Oster-Adjusted (2)	δ^* (3)
Application	0.002	0.004	−0.8
	(0.005)	(0.006)	
N	3,386	3,386	
Written Test	0.0004	0.019***	—†
	(0.007)	(0.007)	
N	3,066	3,066	
Group Debate	−0.003	0.013	0.18
	(0.007)	(0.008)	
N	2,766	2,766	
Interview	−0.004	0.024**	0.18
	(0.011)	(0.011)	
N	2,582	2,582	
Offer	−0.100***	−0.075***	3.8
	(0.014)	(0.015)	
N	2,149	2,149	

Notes: Table H.4 compares OLS estimates (Column 1) to Oster-adjusted estimates (Column 2), with δ^* in Column 3, separately by surname-identifiability sample. Panel A (identifiable surnames, n=619) shows penalties emerging at written tests and persisting through offers. Panel B (non-identifiable surnames, n=3,386) shows near-zero pre-interview effects and a robust offer-stage penalty. All specifications include cohort-year fixed effects and linear controls. The Oster adjustment uses $\delta = 1$ and $R^*_{\max} = 1.3\tilde{R}$. Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. †For stages with near-zero OLS coefficients, δ^* is numerically unstable and not reported.

I A Model Exploiting Plausibly Exogenous Variation in Selectivity Across Jobs to Identify Discrimination

This section formalizes the intuition from the preceding empirical facts. The core challenge is to distinguish discrimination from unobserved ability differences. The interview penalty (Section 6.2) followed by the subsequent promotion advantage (Section 6.3) 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 stylized 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. This selectivity-based approach provides a general framework for detecting discrimination in any multi-stage selection process where post-selection outcomes are observable.

The model is deliberately stylized to clarify the key intuitions and identification assumptions that help interpret the empirical facts through the lens of the model. 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 the qualitative threshold-comparison logic 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 6.4). 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**)**.** This assumption simplifies the exposition by allowing us to treat observed scores as direct measures of ability. Crucially, our main results regarding the identification of discrimination through selectivity variation remain robust even when employers observe only noisy signals of ability.

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:

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

implying each successive round is at least as demanding as the previous one. This comparison is coherent because each observed score $M_{s,i,g}$ is implicitly mapped to a common latent ability metric; all thresholds are therefore expressed on the same scale.

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.

For each screening stage $s \in \{1,2,3,4\}$ and caste group $g \in \{0,1\}$, let $M_{s,i,g,j}$ denote applicant i's (residualized) stage-s score when applying to job j, and let X_i be the observables included in the regressions (GPA, entrance scores, degree/major, experience). There exist distributions $\{F_{s,g}\}$ with supports $\mathcal{S}_{s,g}$ such that, after residualizing on X_i,

$$U_{s,i,g,j} \;\equiv\; M_{s,i,g,j} - m_s(X_i) \;\sim\; F_{s,g} \quad \text{for all eligible jobs } j,$$

with $\text{supp}(F_{s,g}) = \mathcal{S}_{s,g}$ not depending on j. Equivalently, $U_{s,i,g,j} \overset{d}{=} U_{s,i,g,j'}$ for all j, j'.

This assumption is plausible in our setting due to the centralized application portal, near-zero marginal cost of applications, and binding eligibility constraints (Section 5), which minimize caste-based sorting across jobs. Indeed, as shown in Figure 2, average application behavior is similar by caste. However, similar averages could mask differential sorting at the tails—disadvantaged castes might avoid the most selective jobs. Section 6.2 is consistent with this assumption: the absence of a caste penalty even among the most selective jobs suggests that all jobs draw from the same distribution $F_{s,g}$ for each (stage, caste) pair. While common-pool is a maintained assumption, the absence of widening pre-interview caste gaps across selectivity tiers is consistent with it and argues against composition-driven explanations of the interview-stage penalty.

What Assumption 4 does not imply. Note that this assumption does *not impose* equal means or variances across groups. It simply requires that all jobs draw from the same applicant pool—a consulting firm and a tech firm both face the same $F_{s,0}$ and $F_{s,1}$ distributions, rather than attracting systematically different types of advantaged or disadvantaged applicants.

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.

These assumptions lead to the following two corollaries:

Corollary 1 (Exogeneity of cross-job threshold variation)**.** Under Assumptions 1–5, the institutional features documented in Section 5 (centralized portal, near-zero marginal application costs, binding eligibility constraints), and the fact that caste is typically not salient prior to interviews (Section 6.4), the cross-job variation in pre-interview thresholds $\{\tau_{s,j} : s < 4\}$ is orthogonal to group identity *and* to caste-specific unobservables, conditional on observables:

$$\text{For all } j \in \mathcal{J} \text{ and for stages } s < 4: \quad \tau_{s,j} \perp G_i \mid X_i \quad \text{and} \quad \mathbb{E}\big[U_{s,i,g,j} \,\big|\, \tau_{s,j},\, X_i,\, G_i\big] = \mathbb{E}\big[U_{s,i,g,j} \,\big|\, X_i,\, G_i\big],$$

where $U_{s,i,g,j} \equiv M_{s,i,g,j} - m_s(X_i)$ is defined in Assumption 4. Equivalently,

$$\text{Cov}\big(\mathbf{1}\{\text{Disadvantaged}_i\},\, \tau_{s,j} \,\big|\, X_i\big) \;\approx\; 0, \quad s < 4, \qquad \text{and} \qquad \text{Cov}\big(U_{s,i,g,j},\, \tau_{s,j} \,\big|\, X_i,\, G_i\big) \;\approx\; 0, \quad s < 4.$$

These two conditions have distinct sources. The first—$\tau_{s,j} \perp G_i \mid X_i$—is an identifying restriction motivated by the centralized-portal design and the claim that caste is not acted upon before interviews: if firms cannot observe caste at early stages, their thresholds cannot condition on it. The second—mean independence of $U_{s,i,g,j}$ from $\tau_{s,j}$—is a direct implication of Assumption 4's common-pool structure, which posits that the distribution of unobserved ability is job-invariant within (X_i, G_i) cells. Note that mean independence of $U_{s,i,g,j}$ given (X_i, G_i) is sufficient for Propositions 1 and 2 below; full conditional independence is not required.

Corollary 2 (Interview cutoff as a valid instrument). Under the same assumptions and organizational separation between hiring and promotion (HR sets $\tau_{4,j}$; line managers decide promotions later), cross-job variation in the interview cutoff $\tau_{4,j}$ is a valid instrument for the composition of hires: *(i) relevance* $\partial P(\text{Hired}_{j,g} \mid \tau_{4,j})/\partial \tau_{4,j} \neq 0$; *(ii) exclusion* for any job j, group g, and ability index a, the promotion mapping $\pi_{j,g}(a)$ does not depend on $\tau_{4,j}$ once (j, a, g) are fixed, i.e., $\tau_{4,j}$ affects promotions only through selection.

Interpreting the Main Empirical Patterns

We now interpret Sections 6.2 and 6.3 through the lens of the model. As we formalize below, the pattern of no pre-interview gaps but stable post-hire outperformance across selectivity levels is inconsistent with ability-based explanations and aligns with discrimination at interviews.

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. Throughout, $\Delta_{s,j}$ denotes the *gap in conditional pass probabilities* at stage s given that the candidate has reached that stage; earlier-stage selection is baked into $\mathcal{I}_{s-1}$. The proposition is stated for the pooled disadvantaged indicator (SC/ST/OBC combined); subgroup heterogeneity checks are reported separately in Appendix B. The hazard-rate logic still applies to this conditional gap because raising $\tau_{s,j}$ moves both groups toward more selective portions of their stage-specific distributions under the common-pool setup, magnifying any latent right-tail deficit.

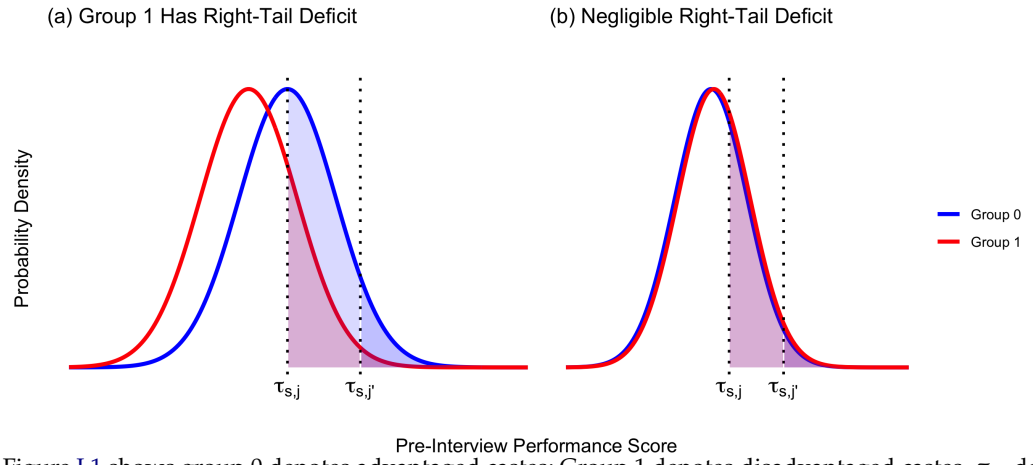
Sketch of Proof. Figure I.1 shows why variation in job selectivity reveals ability differences before interviews. Given Corollary 1, cross-job variation in the pre-interview cutoffs $\tau_{s,j}$ provides *exogenous* shifts that probe progressively deeper into both groups' *common* right tail.

- **Panel (a): Right-tail deficit.** If disadvantaged castes had weaker right-tail performance before interviews, then combining Assumption 5 (light tails/diverging hazard) with

the maintained premise that the disadvantaged distribution is left-shifted relative to the advantaged, 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.** Section 6.2 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 I.1: **Testing for Pre-Interview Performance Differences**



Notes: Figure I.1 shows 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 Section 6.2 show no right-tail deficit before interviews.

□

Ruling out Cumulative Performance Deficits

The absence of early-stage gaps even in highly selective jobs (Section 6.2) rules out not just stage-specific ability differences *but also cumulative performance deficits* across rounds. If disadvantaged castes accumulated small performance disadvantages at each stage that only became decisive at interviews, then by Proposition 1, such cumulative deficits would be most visible as a caste penalty *before interviews* precisely in jobs with high cumulative selectivity—those that impose the most stringent cutoffs at every stage. The key step: a joint pass-rate disadvantage across multiple rounds necessarily implies that at least one stage exhibits a lower conditional pass probability for the disadvantaged group; that stage then satisfies the conditions of Proposition 1 and would reveal a gap at high cutoffs. Intuitively, as any pre-interview cutoff $\tau_{s,j}$ rises, even small per-stage gaps are magnified at that stage; hence a sustained cumulative deficit must surface *before* interviews in highly selective jobs—contrary to Section 6.2. Yet we observe no pre-interview gaps even in these most selective positions—a pattern strongly inconsistent with cumulative performance explanations for the interview penalty.

Do Interviews Reveal Unobserved Ability Differences That Justify Disparities?

Given Proposition 1 and the differential timing of the caste penalty by subsample (Figure 6), 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. This definition assumes that all screening stages, including the interview, provide valid signals of ability, establishing a benchmark against which to test for discrimination. Any strictly monotone transformation preserves the ordering of candidates and therefore the qualitative threshold-comparison logic; I use the average for expositional simplicity. To apply the mean-residual-life result below, I assume a single-index structure: each $M_{s,i,g} = f_s(A_{i,g}) + \varepsilon_{s,i,g}$, where $A_{i,g}$ is one common latent ability draw and $\varepsilon_{s,i,g}$ is small idiosyncratic noise. This simplifying assumption—that stages measure the same underlying ability with minor measurement error—implies $\bar{M}_{i,g}$ and $M_{4,i,g}$ share the same tail behavior. Conditioning on $M_{4,i,g} \geq \tau_{4,j}^{(g)}$ therefore implies a corresponding lower bound on $\bar{M}_{i,g}$ via a monotone mapping $\kappa(\tau_{4,j}^{(g)})$ onto the performance index. For expositional simplicity, I use a tail-index normalization where this mapping is approximately identity at high cutoffs; this is what allows the proposition below to express the limiting gap in cutoff units. This single-index benchmark abstracts from the institutional observation that different stages target somewhat different skills; the identifying content comes from shared tail behavior rather than literal equality of thresholds.

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 $\tau_{4,j}^{q}$ 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 $\tau_{4,j}^{(g)}$ denote the actual group-specific cutoff that caste g faces in job j. Under no discrimination, $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} = \tau_{4,j}^{q}$; under discrimination, $\tau_{4,j}^{(1)} > \tau_{4,j}^{(0)}$. Define the average screening score of job j's hires from caste g as:

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

Because only candidates who cleared earlier rounds reach the interview, conditioning on $M_{4,i,g} \geq \tau_{4,j}^{(g)}$ implicitly conditions on earlier passes; with light tails and increasing thresholds, this is asymptotically equivalent to conditioning on the full intersection of stage cutoffs. Since earlier thresholds are common across castes, this does not affect the caste difference in expected scores.

Define the promotion gap in selectivity tier q as $\Delta^{\text{Prom}}(q) = \bar{Y}_j^{(1)} - \bar{Y}_j^{(0)}$. Terminologically, $\Delta^{\text{Prom}}(q)$ is a gap in *average performance among hires*, used as a promotion proxy because promotion probabilities are increasing in $\bar{M}_{i,g}$ in our setting; a positive $\Delta^{\text{Prom}}(q)$ therefore implies a corresponding promotion-rate advantage under any monotone promotion rule.

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

1. *No discrimination.* If $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} = \tau_{4,j}^{q}$ 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. Given Corollary 2, 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 $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} = \tau$, let $\kappa(\tau)$ denote the induced cutoff on the performance index $\bar{M}$ from the interview-selection event $M_{4,i,g} \geq \tau$. Under the tail-index normalization discussed above, $\kappa(\tau) \approx \tau$ at high cutoffs. Under Assumption 5 (light tails), average performance of hires exceeds this induced cutoff by the Mean Residual Life (MRL), $\mathbb{E}[\bar{M} - \kappa(\tau) \mid \bar{M} \geq \kappa(\tau)] \equiv m(\kappa(\tau))$, so

$$\bar{Y}_j^{(g)} = \kappa(\tau) + m(\kappa(\tau)), \qquad m(\tau) = \frac{\int_\tau^\infty \{1 - F(u)\}\, du}{1 - F(\tau)},$$
$$m'(\tau) = h(\tau)\, m(\tau) - 1, \qquad h(\tau) = \frac{f(\tau)}{1 - F(\tau)}.$$

  With light tails, $h(\tau) \to \infty$ and hence $m(\tau) \downarrow 0$ as $\tau \to \infty$ (full derivation in Appendix Section I.2). Therefore, hired-performance gaps vanish at high selectivity.

- **Panel (b): Discrimination.** If disadvantaged castes face $\tau_{4,j}^{(1)} = \tau_{4,j}^{(0)} + \delta$ with $\delta > 0$, then $\bar{Y}_j^{(1)} - \bar{Y}_j^{(0)}$ converges to a positive constant proportional to this threshold wedge; under the normalization above, that limit is δ. Thus, a positive and stable promotion advantage persists even in the most selective jobs. Because ability-based gaps should vanish at very high selectivity under the maintained assumptions, this pattern is most consistent with disadvantaged castes facing a higher hiring threshold ($\delta > 0$), as illustrated in Figure I.2.

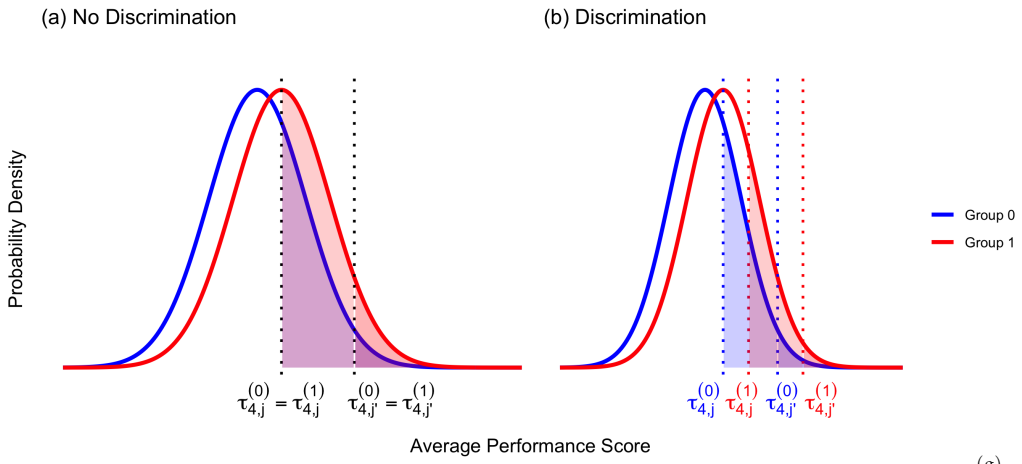
Interpreting Section 6.3 through Assumption 4. As selectivity increases, the promotion rate of hires mechanically converges to that of the marginal hire—this "average-to-marginal" convergence requires only truncation and regularity, not Assumption 4. The assumption adds

interpretation: because $F_{4,g}$ is common across jobs, a given cutoff τ indexes the same margin across firms. Thus, for high selectivity:

$$\underbrace{P_{j,1}(\tau) - P_{j,0}(\tau)}_{\text{promotion gap among hires}} \approx \underbrace{\pi_{j,1}(a(\tau)) - \pi_{j,0}(a(\tau))}_{\text{gap at common margin}},$$

where $\pi_{j,g}(a)$ denotes promotion probability for ability a. If post-hire treatment were equal, the right-hand side would be zero at a common margin. Differences in promotion at high selectivity thus reflect differential treatment (discrimination) rather than different applicant pools. The exclusion restriction—that selectivity $\tau_{4,j}$ affects promotions only through who gets hired—is plausible because HR recruiters set hiring cutoffs while line managers later decide promotions based on observed performance (Section 6.3).

Figure I.2: **Testing for Discrimination**



Notes: Figure I.2 shows 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. □

I.1 Extending the Simple Model of Late-Stage Discrimination with Imperfect Information

In Section I, 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 the *realized* stage-s score to exceed a threshold:

$$M_{s,i,g} \geq \tau_{s,j}, \tag{A1}$$

where $M_{s,i,g}$ is the observed score of candidate i who belongs to group g at round s, and $\tau_{s,j}$ is the job-specific threshold for round s. Let $\mathcal{I}_s \equiv \{\mathcal{I}_{s-1}, M_{s,i,g}\}$ denote the employer's *information set* after observing signals through round s. The decision rule can be written as passing if $\mathbb{E}[A_{i,g} \mid \mathcal{I}_s] \geq \bar{\tau}_{s,j}$. Under symmetric mean-zero noise, this is equivalent to thresholding directly on the realized score $M_{s,i,g} \geq \tau_{s,j}$, because $\mathbb{E}[A_{i,g} \mid \mathcal{I}_s]$ is strictly increasing in $M_{s,i,g}$ given $\mathcal{I}_{s-1}$.

Because only candidates who cleared earlier rounds reach the interview, conditioning on $M_{4,i,g} \geq \tau_{4,j}^{(g)}$ implicitly conditions on earlier passes; with light tails and increasing thresholds, this is asymptotically equivalent to conditioning on the full intersection of stage cutoffs. Since earlier thresholds are common across castes, this does not affect the caste difference in expected scores. (Conditioning on $M_{4,i,g} \geq \tau_{4,j}^{(g)}$ is equivalent to conditioning on "the posterior $\mathbb{E}[A_{i,g} \mid M_4, \mathcal{I}_3] \geq \bar{\tau}_{4,j}^{(g)}$" for a suitably defined threshold $\bar{\tau}$, because for symmetric mean-zero noise the posterior mean of ability is strictly increasing in the realized score.) Throughout, I treat $\eta_{s,i,g}$ as the *innovation* of the stage-s signal—i.e., the new component orthogonal to $\mathcal{I}_{s-1}$—so by construction $\mathbb{E}[\eta_{s,i,g} \mid \mathcal{I}_{s-1}] = 0$ even if raw performance is serially correlated across rounds.

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. (iii) Defining $\eta_{s,i,g}$ as the stage-s innovation relative to $\mathcal{I}_{s-1}$ imposes the martingale-difference property $\mathbb{E}[\eta_{s,i,g} \mid \mathcal{I}_{s-1}] = 0$; this does not require round-wise independence of raw scores or abilities.

Assumptions. Throughout this appendix we maintain the assumptions laid out in Section I 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 Appendix Section I.2.

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 $\mathbb{E}[A_{i,g} \mid \mathcal{I}_s]$ (equivalently on $M_{s,i,g}$ under the symmetry condition above). For high $\tau_{s,j}$, employers require a sufficiently high posterior signal. 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 converges to a positive limit proportional to δ (equal to δ under the normalization in Proposition 2), even in highly selective jobs.

Hence, persistent promotion gaps at high selectivity are difficult to reconcile with an unobserved-ability explanation and are more consistent with discrimination under the maintained assumptions.

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. (Recall that for any light-tailed distribution, $\mathbb{E}[A - \tau \mid A \geq \tau] \equiv m(\tau) \to 0$ as $\tau \to \infty$; see Section I.2.) Under imperfect information, selection at stage 4 is the event $M_{4,i,g} = A_{i,g} + \eta_{4,i,g} \geq \tau_{4,j}^{(g)}$. 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 the relevant decomposition for group 1 is:

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

Defining $\eta_{4,i,g}$ as the round-4 innovation implies $\mathbb{E}[\eta_{4,i,g} \mid \mathcal{I}_3] = 0$, but selection depends on realized $M_{4,i,g}$, so the relevant object is the selected mean above. Under the high-cutoff light-tail conditions in Proposition 5, this selected overshoot term is asymptotically negligible. Hence the limiting promotion-gap logic is unchanged: under equal thresholds gaps vanish, while under a threshold wedge they converge to a positive limit (equal to δ under the normalization in Proposition 2).

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.

□

I.2 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 light-tailed distributions with rising hazards commonly used in practice (e.g. Normal, Logistic, and Weibull with shape parameter > 1), but excludes constant-hazard cases such as Exponential.

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

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

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. Since $m(\tau) = \frac{\int_\tau^\infty S(u)\, du}{S(\tau)}$, L'Hôpital's rule gives

$$\lim_{\tau \to \infty} m(\tau) = \lim_{\tau \to \infty} \frac{\int_\tau^\infty S(u)\, du}{S(\tau)} = \lim_{\tau \to \infty} \frac{-S(\tau)}{-f(\tau)} = \lim_{\tau \to \infty} \frac{S(\tau)}{f(\tau)} = \lim_{\tau \to \infty} \frac{1}{h(\tau)} = 0$$

under Assumption A (finite mean, $h(\tau) \to \infty$).

This approach avoids any unstated monotonicity assumption and relies only on the explicitly stated conditions in Assumption A. □

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. Constant-hazard cases (e.g., the Exponential distribution) also fail this assumption and imply a non-vanishing mean residual life. In our context, performance scores are effectively bounded and therefore satisfy the needed tail regularity, so the limit result holds for the objects used in Proposition 2.

J LLM-Based Transcript Analysis Methodology

This appendix provides complete documentation of the LLM-based transcript analysis methodology, including the unified classification scheme, prompts, and robustness checks.

J.1 Key Design Principle: Unified Taxonomy

The central methodological innovation is applying the *identical classification schema* to both HR interviews and group debates. The classifier records 21 raw topic dimensions—10 policy/societal topics and 11 biographical topics (including college education)—regardless of format. For main comparative interpretation, we exclude college education and use a 20-topic analytic set (10 policy + 10 biographical markers). This ensures the comparison is fair and any differences reflect genuine content differences rather than classification artifacts.

Why does this matter? Consider the alternative: if we designed one taxonomy for HR interviews (focused on biographical markers) and a different taxonomy for GDs (focused on policy topics), any observed differences could be dismissed as artifacts of differential measurement. By applying the *same* 20 categories to *both* formats, we impose no prior assumptions about which format "should" probe which topics. The finding that HR interviews concentrate on biographical markers while GDs concentrate on policy topics emerges from the data, not from our classification design. This also makes the comparison conservative: GDs *could* in principle probe biographical content (the taxonomy allows it), but they empirically do not.

Critical Distinction. We distinguish between biographical content *as the topic* versus personal experience used to discuss societal topics:

- **Biographical** = The conversation extracts personal facts about the candidate AS THE MAIN TOPIC. Example: "Where are you from?" $\rightarrow$ "I'm from Jaipur" (hometown is the topic).
- **Not Biographical** = Personal experience used as context for discussing a SOCIETAL topic. Example: "As someone from a rural area, I think infrastructure investment is crucial" (economic policy is the topic, not the speaker's origins).

J.2 Data Collection

We systematically collected mock interview transcripts from YouTube, the dominant platform for interview preparation content in India. Our search strategy combined general terms ("HR interview mock," "personal interview mock India," "group debate mock," "GD practice") with company-specific queries targeting the multinationals in our sample (e.g., "TCS interview," "Bosch GD"). To ensure we analyze content that candidates actually consume, we retained only videos with at least 1,000 views; the most-watched videos have view counts in the millions, confirming substantial engagement with this preparation material.

Transcripts were extracted using commercial transcription services—primarily Notta AI and Reduct Video—selected for their ability to handle mixed Hindi-English ("Hinglish") dialogue common in Indian interview settings. YouTube's automatic captions served as a fallback where commercial transcription was unavailable. All transcripts underwent manual review for transcription errors. To prevent counting the same interview multiple times when

re-uploaded across channels, we identified and removed duplicates using content-based hashing (MD5 of normalized text).

The final corpus comprises 390 transcripts: 202 group debate transcripts and 188 HR interview transcripts. After filtering to live mock sessions (excluding tips and meta-discussion), we retain 274 transcripts for analysis: 148 live group debates (622,322 words; average 4,204 words per transcript) and 126 live HR interviews (247,956 words; average 1,967 words per transcript), totaling approximately 870,000 words. Relative to earlier working-draft live counts, this reflects expanded group-debate collection plus stricter live-versus-tips filtering and exclusion of personally conducted HR-interview files from analysis. Many videos explicitly target preparation for the multinationals studied in this paper, helping align the analyzed dialogue with actual campus placement processes.

Filtering to Live Transcripts

Many YouTube preparation videos contain tips, advice, or meta-discussion rather than actual mock interviews. After classifying each transcript by type (see Section J.3 below), we retain only "live" mock sessions—actual simulated interviews with question-and-answer exchanges or group debates with substantive participant interaction. This filtering yields 274 live transcripts for analysis: 148 live group debates and 126 live HR interviews. All results reported in the main text derive from these live transcripts to ensure we compare actual interview content rather than coaching advice.

J.3 Classification Prompts

We implemented the classification using GPT-4o-mini via the OpenAI API. We chose this model for its balance of capability and cost-efficiency given our large corpus; preliminary tests showed classification quality comparable to GPT-4 at substantially lower cost. We used structured JSON outputs to ensure consistent parsing and set temperature to 0.1 to minimize response variance across transcripts. The prompt design reflects several critical methodological choices, which we explain before presenting the complete prompts.

Unified 20-Category Schema. We apply an *identical* classification schema to both HR interviews and group debates. The schema comprises 10 policy/societal topics (economic policy, education policy, technology and AI, gender and women's rights, environment, governance and democracy, media and social media, society and culture, food policy, and other policy topics) and 10 biographical topics (hometown/origins, hobbies and interests, high school attended, family structure, native place versus current city, parents' occupation, school type, coaching/tuition history, medium of schooling, and school board). (We exclude college/university education because all candidates already attend an elite institution; see Figure 5 notes.) By evaluating both formats against the same dimensions, we ensure that any observed differences reflect genuine content variation rather than differential classification criteria. We do not assume *a priori* which format "should" probe which topics—the model applies identical standards to both.

Biographical Content as Topic versus Context. The most critical design choice is the explicit distinction between biographical content *as the topic* versus personal experience used as *context* for policy discussion. This distinction prevents false positives in group debates. Consider the contrast: when an interviewer asks "Where are you from?" and the candidate

responds "I'm from Jaipur," hometown is the topic being extracted—this is biographical. But when a GD participant says "As someone from a rural area, I think infrastructure investment is crucial," rural origins serve as rhetorical context for an economic policy argument—this is *not* biographical. Group debate participants frequently invoke personal experiences to lend credibility to policy arguments; without this rule, such rhetorical moves would be falsely classified as biographical extraction, artificially inflating the biographical content of GDs and biasing the comparison. The prompt provides explicit examples of both cases to guide the model's classification.

Purely Descriptive Classification. The prompt is purely descriptive: it identifies what topics are covered, not how well they are covered. The model does not evaluate response quality, interview performance, or candidate competence. It does not assess whether answers are "good" or "bad," whether candidates appear nervous or confident, or whether responses are substantively correct. Nor does it weight biographical topics differently for HR versus GD transcripts—a question about hometown is classified identically regardless of which format it appears in. The goal is objective content classification, not evaluative judgment.

What the Prompt Does Not Do. Several potential sources of bias are explicitly avoided. The prompt does not: (1) make assumptions about which format "should" contain which content; (2) treat personal anecdotes in GDs as biographical unless personal facts are the explicit topic of conversation; (3) penalize transcripts for code-switching or non-native English; (4) consider the order or timing of topics within transcripts; or (5) attempt to identify individual speakers' caste, class, or demographic characteristics. The prompt focuses solely on content classification at the transcript level.

Output Structure. The prompt returns structured JSON with the following fields: transcript type (live interview, live group debate, tips, or meta-content), high-level content focus (primarily about the speaker versus primarily about society, with a "mixed" option for genuinely balanced transcripts), confidence level, estimated percentage of dominant content, and detailed topic analysis across all 20 dimensions. Each topic field includes both a boolean indicator of whether the topic was discussed and a free-text detail field capturing specifics. This structured output eliminates parsing ambiguity and enables systematic aggregation across the corpus.

The complete prompts follow below.

System Prompt

Prompt chunk 1: objective and scope. This chunk forces transcript-level content identification without evaluative scoring.

```
You are analyzing interview/discussion transcripts to identify what topics are
covered.

Your task is to objectively identify:  1.  HIGH-LEVEL FOCUS: Is this primarily
about the speaker/candidate personally, or about society/policy?  2.  POLICY
TOPICS: What societal/policy topics are discussed (if any)?  3.  BIOGRAPHICAL
TOPICS: What personal facts about the candidate are extracted (if any)?
```

Prompt chunk 2: guardrail on biographical extraction. This chunk is the key false-positive guardrail used in both HR and GD transcripts.

```
CRITICAL DISTINCTION FOR BIOGRAPHICAL TOPICS:

BIOGRAPHICAL = The conversation EXTRACTS personal facts about the candidate
AS THE MAIN TOPIC. - "Where are you from?" → "I'm from Jaipur" = BIOGRAPHICAL
(hometown is the topic) - "What does your father do?" → "He's a farmer" =
BIOGRAPHICAL (parents' occupation is the topic) - "Tell me about your family" →
"I have two siblings" = BIOGRAPHICAL (family structure is the topic)

NOT BIOGRAPHICAL = Personal experience used to discuss a SOCIETAL topic.  - "In
my experience at college, reservation policy needs reform" = NOT biographical
(policy is the topic) - "When I worked at Tata, I saw how AI is transforming
manufacturing" = NOT biographical (technology is the topic) - "As someone from
a rural area, I think infrastructure investment is crucial" = NOT biographical
(economic policy is the topic)
```

Prompt chunk 3: consistency requirement.

```
The key question:  Is the TOPIC the person's biography, or is personal
experience just CONTEXT for discussing society/policy?

Apply the SAME analysis to all transcripts regardless of format.  Be thorough
and objective.
```

Unified Classification Prompt (Same for HR and GD)

Prompt chunk 1: transcript-type and high-level labels.

```
Analyze this transcript and identify ALL topics covered.

IMPORTANT: Apply the same analysis regardless of transcript type (GD or HR
interview).

STEP 1 - TRANSCRIPT TYPE: - live_gd:  A live group debate with participants
debating - live_hr:  A live HR/personal interview with Q&A - tips:
Coaching/tips content about interviews or GDs - meta:  Discussion about
interviews/GDs without actual content - other:  Something else

STEP 2 - HIGH-LEVEL CLASSIFICATION: Based on the DOMINANT content (80%+
of transcript), is this:  - about_speaker:  Primarily extracts/discusses
personal information about the candidate - about_society:  Primarily discusses
societal/policy topics - mixed:  Genuinely balanced (40-60% each) - use
sparingly - If one side is 60-80%, classify to that dominant side (about_speaker
or about_society) and mark confidence as medium.
```

Prompt chunk 2: topic taxonomy and extraction rule.

```
STEP 3 - DETAILED TOPIC ANALYSIS: Identify which specific topics are covered.
Use the SAME criteria for both GD and HR.

CRITICAL: For biographical topics, mark TRUE only if personal info is EXTRACTED
AS THE TOPIC. Do NOT mark TRUE if someone merely mentions personal experience to
argue about policy.
```

Prompt chunk 3: structured output schema (part A).

```
Return ONLY a valid JSON object with these top-level keys:  - transcript_type:
live_gd | live_hr | tips | meta | other - content_focus:  about_speaker |
about_society | mixed - content_focus_confidence:  high | medium | low -
dominant_content_pct:  estimated share statement - policy_topics:  nested object
with {discussed, detail} by topic - biographical_topics:  nested object with
{discussed, detail} by topic
```

Prompt chunk 4: structured output schema (part B).

```
Policy-topic keys:  economic_policy, education_policy, technology_ai,
gender_women, environment, governance_democracy, media_social_media,
society_culture, food_policy, other_policy.

Biographical-topic keys:  hometown_origins, hobbies_interests, high_school,
family_structure, native_vs_current, parents_occupation, school_type,
coaching_tuition, medium_of_schooling, school_board, college_education.

For reported ‘‘biographical markers’’ figures/tables, college_education is
excluded, leaving a 10-marker set used in the main interpretation.
```

J.4 Classification Process

The classification proceeds in two steps using the unified prompt above:

Step 1: High-Level Content Focus. Each transcript is classified by dominant content—"about_speaker" when personal extraction is dominant, "about_society" when policy/societal discussion is dominant, and "mixed" only for genuinely balanced 40–60 splits. Cases in the 60–80 range are assigned to the dominant side with medium confidence.

Step 2: Detailed Topic Identification. For each transcript, we identify which of the 20 specific topics are covered. The same 10 policy topics and 10 biographical topics are evaluated for both HR interviews and group debates.

J.5 Results Summary

Results are based on the 274 live transcripts (126 HR interviews, 148 group debates). The corpus spans real mock-preparation content that candidates actively use before placement rounds, not synthetic prompts generated for this study. It therefore captures recurring

question formats, topic sequencing, and interviewer-candidate exchanges at scale. The HR-versus-GD contrast is informative because both are communication-intensive rounds, but only one systematically elicits personal biography as the topic of evaluation.

High-Level Content Classification

Table J.1: Content Focus Classification

Content Focus	HR Interviews (n=126)	Group Debates (n=148)
About Speaker (Personal/Biographical)	89.7%	0.7%
About Society (Policy/Debate)	3.2%	99.3%
Mixed	7.1%	0.0%

Notes: Table J.1 shows high-level content classification of 274 live transcripts from YouTube mock interview preparation videos. Each transcript is classified based on its dominant content (80%+ of transcript): "About Speaker" if the conversation primarily extracts personal/biographical information about the candidate; "About Society" if it primarily discusses policy/societal topics; "Mixed" if genuinely balanced (40–60% each). Classification performed using GPT-4o-mini with the unified prompt described in Section J. Cross-LLM validation with Claude Sonnet 4 and `gpt-5.2` is reported below.

Detailed Topic Analysis: All 21 Dimensions

Table J.2 and Table J.3 report the full topic-level breakdown from the unified taxonomy. Table J.2 summarizes policy-content prevalence (for example, economic policy, governance, and technology) across HR interviews and group debates. Table J.3 then reports extraction rates for biographical markers such as hometown, family background, and schooling trajectories. Reading the two tables jointly clarifies that the format contrast is structural: group debates concentrate on policy debate, while personal interviews disproportionately extract personal background information.

Table J.2: Policy Topics (10 Categories)

Policy Topic	HR Interviews (n=126)	Group Debates (n=148)
Economic Policy	4.8%	55.4%
Education Policy	2.4%	33.8%
Technology & AI	25.4%	29.1%
Gender & Women's Rights	4.8%	18.2%
Environment	8.7%	10.1%
Governance & Democracy	1.6%	14.9%
Media & Social Media	2.4%	13.5%
Society & Culture	8.7%	88.5%
Food Policy	0.8%	6.8%
Other Policy Topics	1.6%	1.4%

Notes: Table J.2 shows percentage of transcripts in which each policy topic was discussed. Sample comprises 274 live transcripts (126 HR interviews, 148 group debates) from YouTube mock interview preparation videos. A topic is marked as "discussed" if the LLM classifier detected substantive discussion of that topic as a main theme (not merely passing mention). Classification performed using GPT-4o-mini with the unified 20-category taxonomy applied identically to both interview formats.

Table J.3: Biographical Topics (10 Categories)

Biographical Topic	**HR Interviews** (n=126)	**Group Debates** (n=148)
Hometown/Geographic Origins	69.8%	7.4%
Hobbies & Interests	67.5%	0.7%
High School Attended	48.4%	4.7%
Family Structure	27.0%	2.7%
Parents/Family Occupation	21.4%	0.7%
Native Place vs Current City	11.1%	1.4%
School Type (Govt/Private/Convent)	7.9%	1.4%
Coaching/Tuition	3.2%	0.0%
Medium of Schooling	2.4%	0.7%
School Board (CBSE/ICSE/State)	1.6%	0.0%

Notes: Table J.3 shows percentage of transcripts in which each biographical topic was probed. Sample comprises 274 live transcripts (126 HR interviews, 148 group debates) from YouTube mock interview preparation videos. A biographical topic is marked as "discussed" only when personal information is *extracted as the main topic* (e.g., "Where are you from?" → "I'm from Jaipur"), not when personal experience is used as context for policy discussion (e.g., "As someone from a rural area, I think infrastructure investment is crucial"). This distinction is critical: the prompt explicitly instructs the classifier to distinguish biographical extraction from rhetorical use of personal experience. Classification performed using GPT-4o-mini with the unified 20-category taxonomy.

J.6 Sampling-Based Sensitivity Analysis

To ensure our findings are not artifacts of the classification methodology, we tested robustness across three sampling strategies that differ in how much of each transcript the LLM observes.

The *full transcript* strategy, which serves as our primary analysis, provides the complete transcript to the model in a single API call. For very long transcripts exceeding 12,000 characters, we retain the first and last 6,000 characters, preserving opening questions and closing exchanges while managing token limits. This approach maximizes context, allowing the model to observe the complete conversational flow.

The *paragraph-by-paragraph* strategy classifies each paragraph independently using a simplified prompt, then aggregates results at the transcript level. Transcripts are split on double newlines or speaker changes, with each paragraph limited to 2,000 characters and a maximum of 20 paragraphs per transcript. For content focus, we apply majority voting with thresholds: a transcript is classified as "about the speaker" if at least 60% of its paragraphs receive that label, "about society" if at least 60% receive that label, and "mixed" if neither threshold is met. For topic identification, we apply union logic: if any paragraph mentions a topic, the transcript is marked as containing that topic. Transcripts where more than 30% of paragraphs are classified as "unclear" receive a low-confidence flag.

The *random sample* strategy selects 30% of paragraphs (minimum 3, maximum 20) using a fixed random seed for reproducibility, then applies the same aggregation logic as the paragraph-by-paragraph approach. This provides a cost-effective robustness check and tests whether findings hold even with partial transcript information.

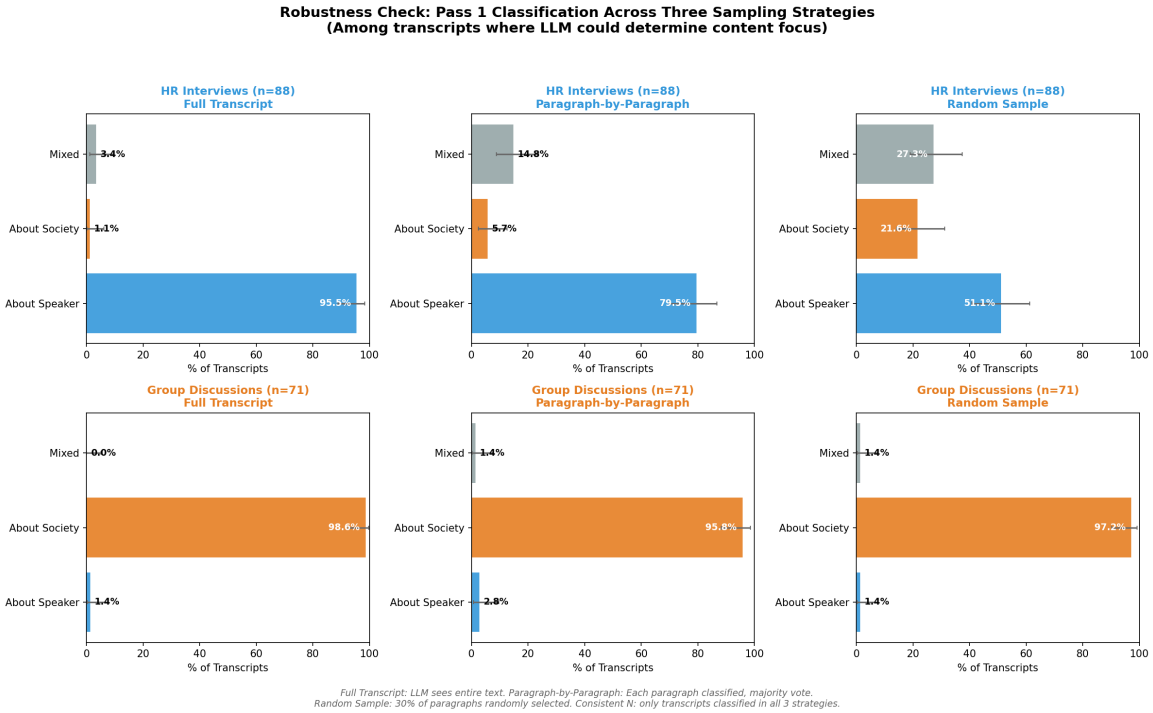
If all three strategies yield similar HR-GD differences, the finding is robust to methodological choices. For the robustness figures, we use consistent sample sizes across strategies, including only transcripts successfully classified in all three approaches. This ensures an apples-to-apples comparison: the same transcripts contribute to each panel.

Crucially, the paragraph-by-paragraph and random sampling strategies *necessarily* produce less stark differences than full-transcript analysis because the LLM has far less context. When classifying a single paragraph or a random 30% sample, the model cannot observe the overall conversational arc—biographical questions that anchor an interview may fall outside the sampled excerpt. This reduced context attenuates measured differences, as we would expect. The key finding is that even under these informationally impoverished conditions, the qualitative story remains the same: HR interviews consistently show higher prevalence of biographical probing than GDs across all sampling strategies. The persistence of this pattern despite reduced context strengthens confidence that the underlying content difference is real.

Results Across Sampling Strategies

Figure J.1 shows the content focus classification across all three strategies. The core finding—HR interviews focus on the speaker while GDs focus on society—is robust across strategies. The 89.7% HR "about speaker" rate refers to the unrestricted full-transcript estimate in Table J.1; Figure J.1 instead plots a restricted consistent-N subset (only transcripts successfully classified under all three strategies), so the full-strategy bar level in the figure can differ. Lower rates in paragraph and random strategies reflect stricter aggregation rules: only transcripts where at least 60% of paragraphs are classified as "about speaker" receive that label at the transcript level, and random sampling may miss biographical questions that appear in specific parts of interviews. GDs are still classified as "about society" at very high rates across all three strategies, confirming the robustness of the core finding.

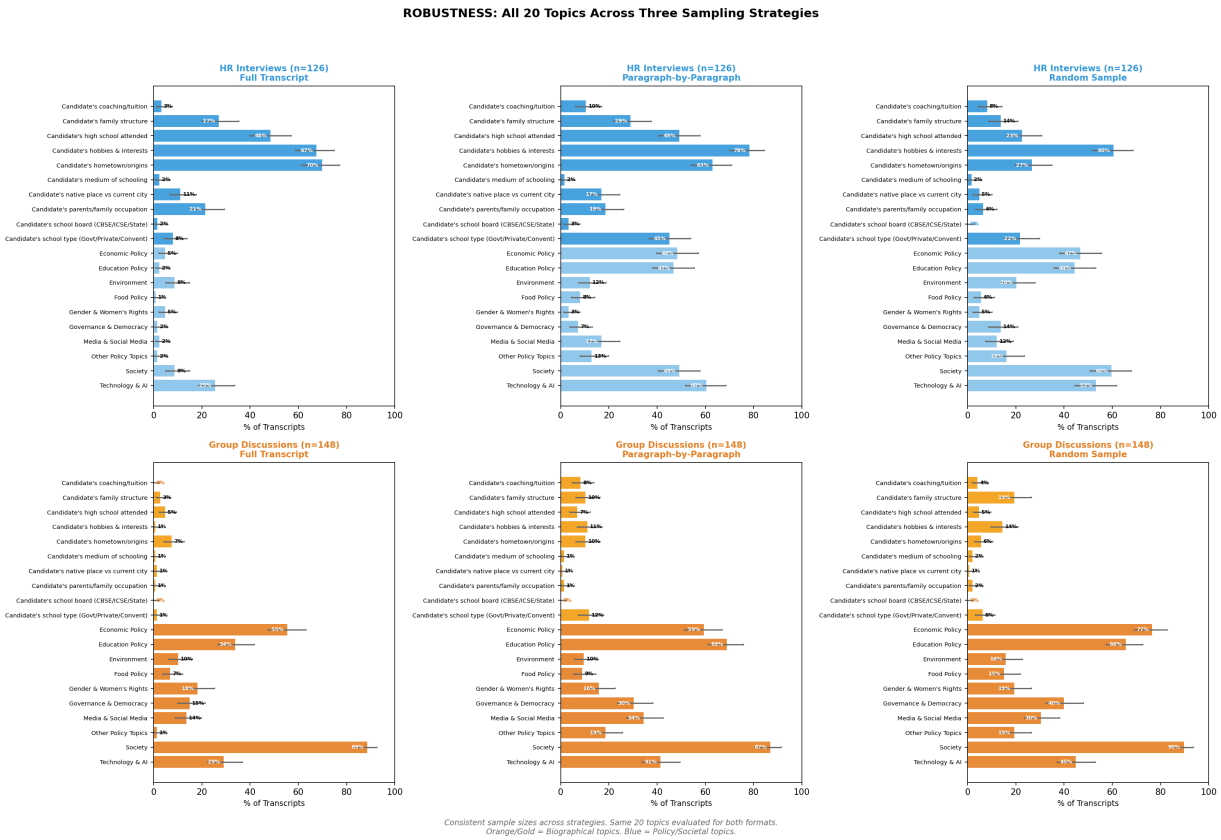
Figure J.1: **Robustness: Content Focus Classification Across Sampling Strategies.**



Notes: Figure J.1 shows high-level content focus classification across three sampling strategies. "Full" uses the complete transcript (truncated to 12,000 characters if necessary); "Paragraph" classifies each paragraph independently and applies majority voting (60% threshold); "Random" samples 30% of paragraphs. Sample includes only transcripts successfully classified in all three approaches for apples-to-apples comparison. Error bars show 95% Wilson confidence intervals. Classification performed using GPT-4o-mini with the unified prompt described in Section J. Lower "about speaker" rates in alternative strategies are expected: with less context, the LLM cannot observe the full conversational arc, so classifications appear more mixed. The qualitative pattern—HR interviews focus on the speaker, GDs focus on society—persists across all strategies.

Figure J.2 shows the detailed topic prevalence across all three strategies for all 20 dimensions. The relative differences between HR and GD are preserved: HR interviews consistently show higher prevalence of biographical topics while GDs show higher prevalence of policy topics, regardless of sampling strategy.

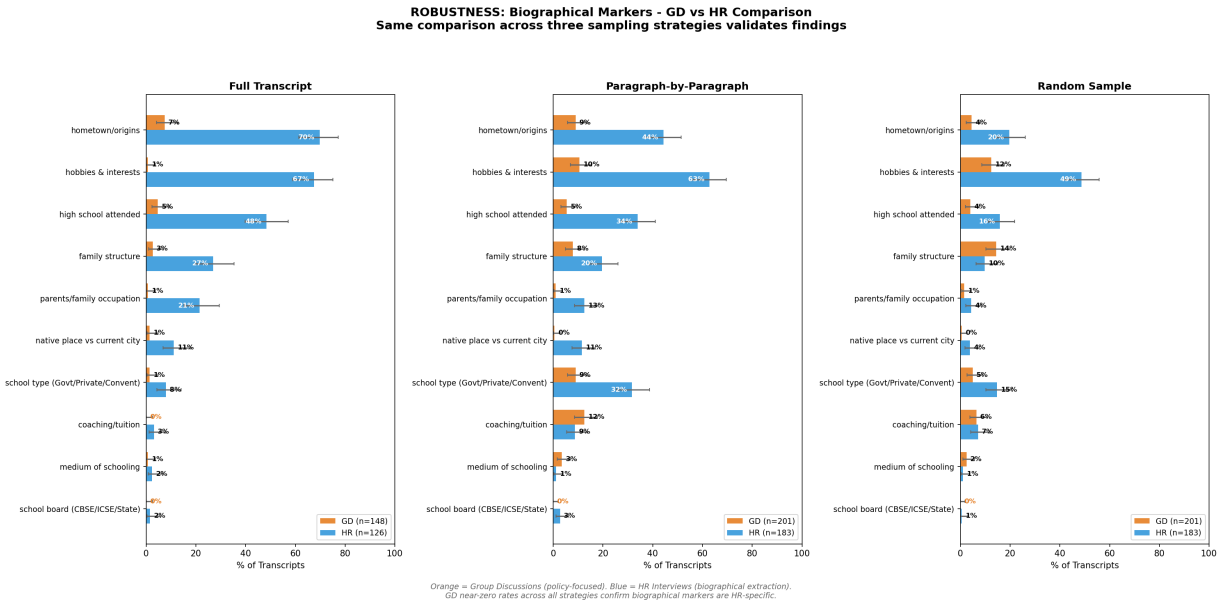
Figure J.2: **Robustness: All 20 Topic Dimensions Across Sampling Strategies.**



Notes: Figure J.2 shows prevalence of all 20 topic dimensions (10 policy, 10 biographical) across three sampling strategies. Each bar shows the percentage of transcripts in which that topic was identified. Colors denote topic family within each subplot: policy topics are shown in blue shades and biographical topics in warm shades, with exact hue varying by format (GD versus HR); read family assignment from each subplot legend. "Full" uses the complete transcript; "Paragraph" classifies each paragraph independently; "Random" samples 30% of paragraphs. For paragraph and random strategies, a topic is marked present if *any* paragraph mentions it (union logic). Sample includes only transcripts successfully classified in all three approaches. Error bars show 95% Wilson confidence intervals. Classification performed using GPT-4o-mini. Because the alternative strategies provide less context, detected prevalence rates are generally lower—but the relative ordering (HR higher on biographical topics, GD higher on policy topics) remains consistent, demonstrating that the qualitative pattern is not an artifact of methodology.

Figure J.3 zooms in on the 10 biographical markers (excluding college education), showing robustness across the three sampling strategies. The large HR-GD differences for markers like hometown/origins, hobbies, and parents' occupation are preserved across strategies.

Figure J.3: **Robustness: Biographical Markers Across Sampling Strategies.**



Notes: Figure J.3 shows prevalence of the 10 biographical marker dimensions (excluding college education) across three sampling strategies, focusing on the topics most relevant to caste identification. Each bar shows the percentage of transcripts in which that biographical topic was probed. A topic is marked as "discussed" only when personal information is *extracted as the main topic*, not when personal experience is used as context for policy discussion. "Full" uses the complete transcript; "Paragraph" classifies each paragraph independently; "Random" samples 30% of paragraphs. Sample includes only transcripts successfully classified in all three approaches. Error bars show 95% Wilson confidence intervals. Classification performed using GPT-4o-mini. With reduced context, the LLM detects fewer biographical markers—this attenuation is expected since biographical questions may fall outside the sampled excerpts. The key finding is that HR-GD differences persist qualitatively: HR interviews consistently show higher biographical probing than GDs across all strategies, even when the model observes only fragmentary portions of each conversation.

Cross-LLM Robustness: Claude and GPT-5.2 Validation

A potential concern is that our findings might reflect artifacts of GPT-4o-mini's classification behavior rather than genuine content differences. To address this, we replicated the entire analysis using Anthropic's Claude Sonnet 4 and OpenAI's `gpt-5.2` (run in February 2026), independently developed models with different architectures and training pipelines. All three models received identical prompts and applied the same unified taxonomy.

Table J.4 presents the comparison. The core finding is robust across models. For high-level content focus, GPT-4o-mini classifies 89.7% of HR interviews as "about speaker," Claude classifies 93.5%, and GPT-5.2 classifies 94.9%. All three models classify essentially all group debates as "about society" (99.3%, 100.0%, and 99.1%, respectively). For biographical topics, all three models find high HR prevalence for hometown/origins (69.8%, 75.7%, 74.5%) and hobbies (67.5%, 72.9%, 70.4%), while GD prevalence stays low (e.g., hometown/origins at 7.4%, 10.2%, and 7.5%). Minor differences in absolute percentages do not affect the qualitative conclusion: HR interviews systematically extract biographical information that group debates do not.

Table J.4: **Cross-LLM Validation: GPT-4o-mini vs. Claude Sonnet 4 vs. GPT-5.2**

	GPT-4o-mini		Claude Sonnet 4		GPT-5.2	
	HR	GD	HR	GD	HR	GD
High-Level Focus						
About Speaker	89.7%	0.7%	93.5%	0.0%	94.9%	0.0%
About Society	3.2%	99.3%	4.7%	100.0%	3.1%	99.1%
Biographical Topics						
Hometown/origins	69.8%	7.4%	75.7%	10.2%	74.5%	7.5%
Hobbies & interests	67.5%	0.7%	72.9%	0.8%	70.4%	0.9%
High school attended	48.4%	4.7%	46.7%	1.6%	48.0%	0.0%
Family structure	27.0%	2.7%	30.8%	0.8%	26.5%	0.9%
Parents' occupation	21.4%	0.7%	23.4%	0.0%	25.5%	0.0%
Sample size (live transcripts)	126	148	107	127	98	107

Notes: Table J.4 shows classification results using identical prompts and unified taxonomy across GPT-4o-mini, Claude Sonnet 4, and `gpt-5.2`. Sample size differences reflect how each model classified ambiguous transcripts (tips vs. live); the comparison uses transcripts classified as "live" by each model. This tests the joint stability of both the live/tips filtering step and the topic classification step—if either step were model-dependent in consequential ways, the core pattern would diverge. The fact that all three models yield similar biographical-vs-policy differences despite operating on somewhat different "live" samples strengthens the robustness finding. Minor percentage differences do not affect the qualitative conclusion: HR interviews systematically probe biographical content while group debates focus on policy debates.

To complement side-by-side prevalence comparisons, Table J.5 reports formal concordance on dominant-content labels for the overlap of transcripts classified as live by each model.

Table J.5: **Cross-LLM Dominant-Content Concordance on Overlapping Live Transcripts**

Model Pair	Overlap N	Agreement (%)	Cohen's κ
GPT-4o-mini vs. Claude Sonnet 4	234	96.2	0.925
GPT-4o-mini vs. GPT-5.2	205	98.1	0.962
Claude Sonnet 4 vs. GPT-5.2	199	98.5	0.970
All three models (exact-match agreement)	199	97.0	–

Notes: This table reports concordance for dominant-content classification (about speaker / about society / mixed), computed on transcript overlaps keyed by filename and source type after applying the same exclusions used in the analysis pipeline (including removal of personally conducted HR-interview files and retention of live transcripts only). "Agreement" is raw percent agreement; κ is Cohen's kappa. Metrics are generated from `scripts/compute_cross_llm_agreement.py`.

Cross-LLM Visual Replications

To keep the visual robustness checks readable, this subsection reports full-size model-specific replications rather than horizontally stacked cross-model composites. Table J.4 remains the harmonized quantitative comparison; the figures below are intended to let readers inspect each model's internal pattern directly.

Claude Sonnet 4 Robustness Visuals

Figures J.5, J.6, and J.7 show the Claude Sonnet 4 replications for content focus, topic dimensions, and biographical markers across full-transcript, paragraph, and random sampling strategies. These are computed on Claude's own "live" subset after model-specific tips-versus-live filtering.

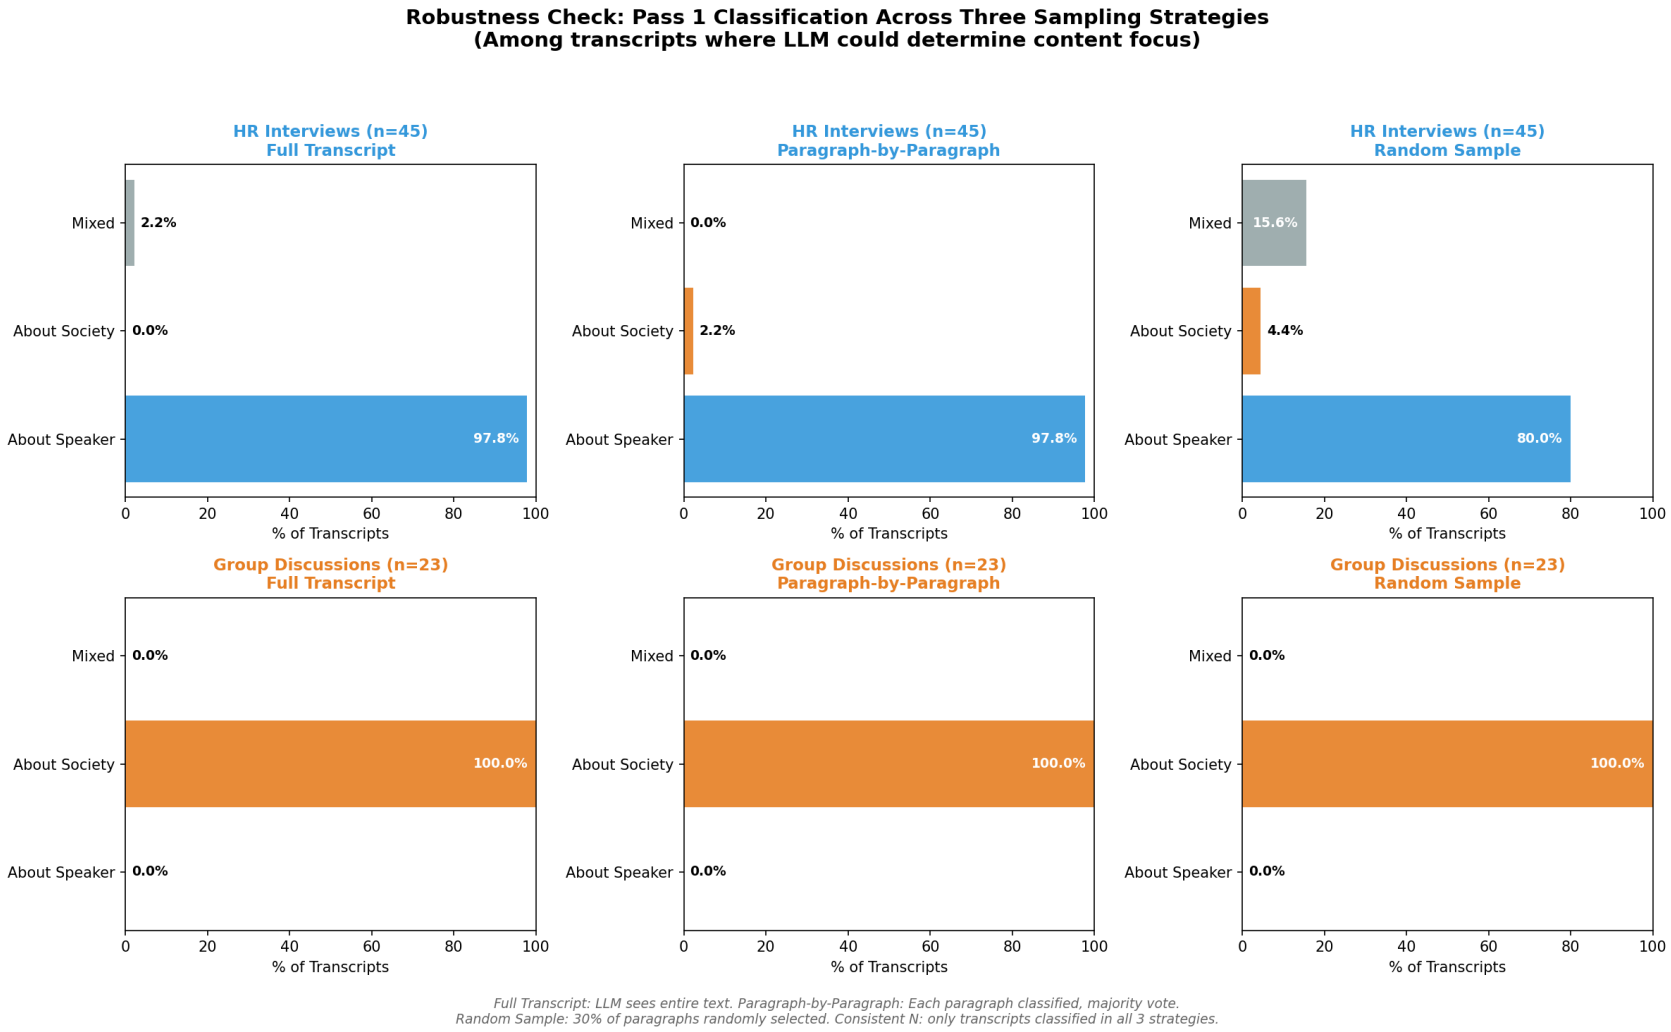


Figure J.5: **Claude Sonnet 4: Robustness Content Focus Across Sampling Strategies**

Notes: Figure J.5 reports the content-focus robustness check for Claude Sonnet 4 under the three sampling strategies used in Appendix J (full transcript, paragraph-by-paragraph, random sample). Bars report the share of transcripts classified as about speaker, about society, or mixed, separately for HR interviews and group debates. The figure uses Claude's model-specific "live" subset and strategy-consistent classified transcripts. Cross-model percentage comparisons are reported in Table J.4.

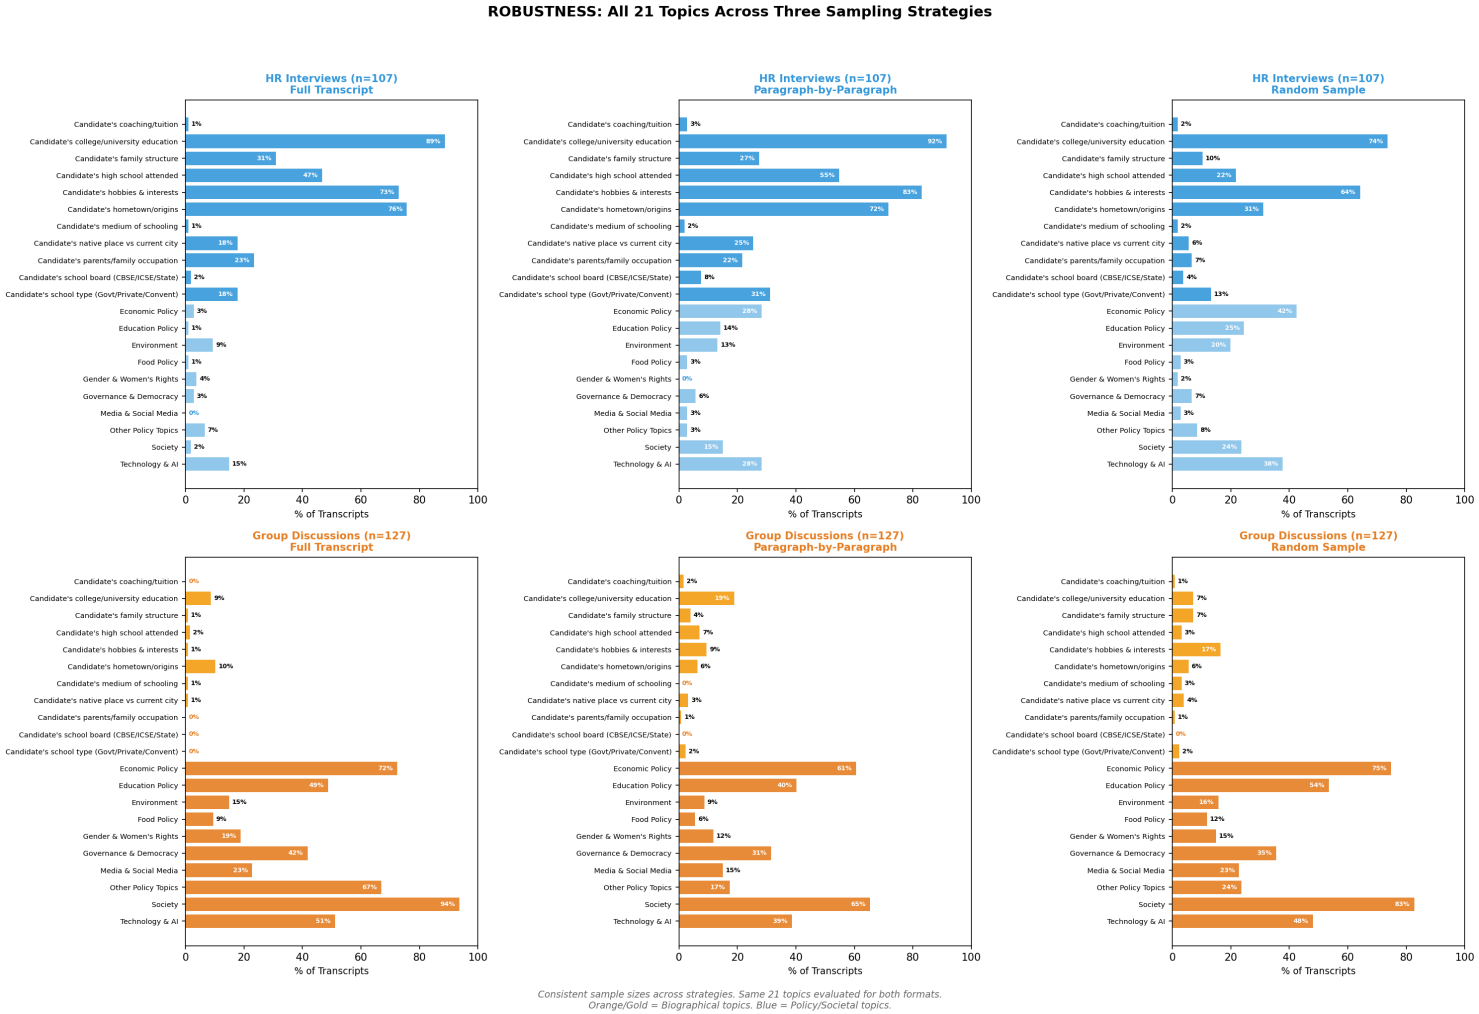


Figure J.6: **Claude Sonnet 4: Robustness Topic Dimensions Across Sampling Strategies**

Notes: Figure J.6 reports the 20-topic robustness visualization for Claude Sonnet 4 (10 policy topics and 10 biographical topics) across full, paragraph, and random sampling strategies. Topic coding uses the same unified taxonomy and prompt family as the main Appendix J analysis. The objective is within-model pattern stability across sampling strategies; exact cross-model percentages are summarized in Table J.4.

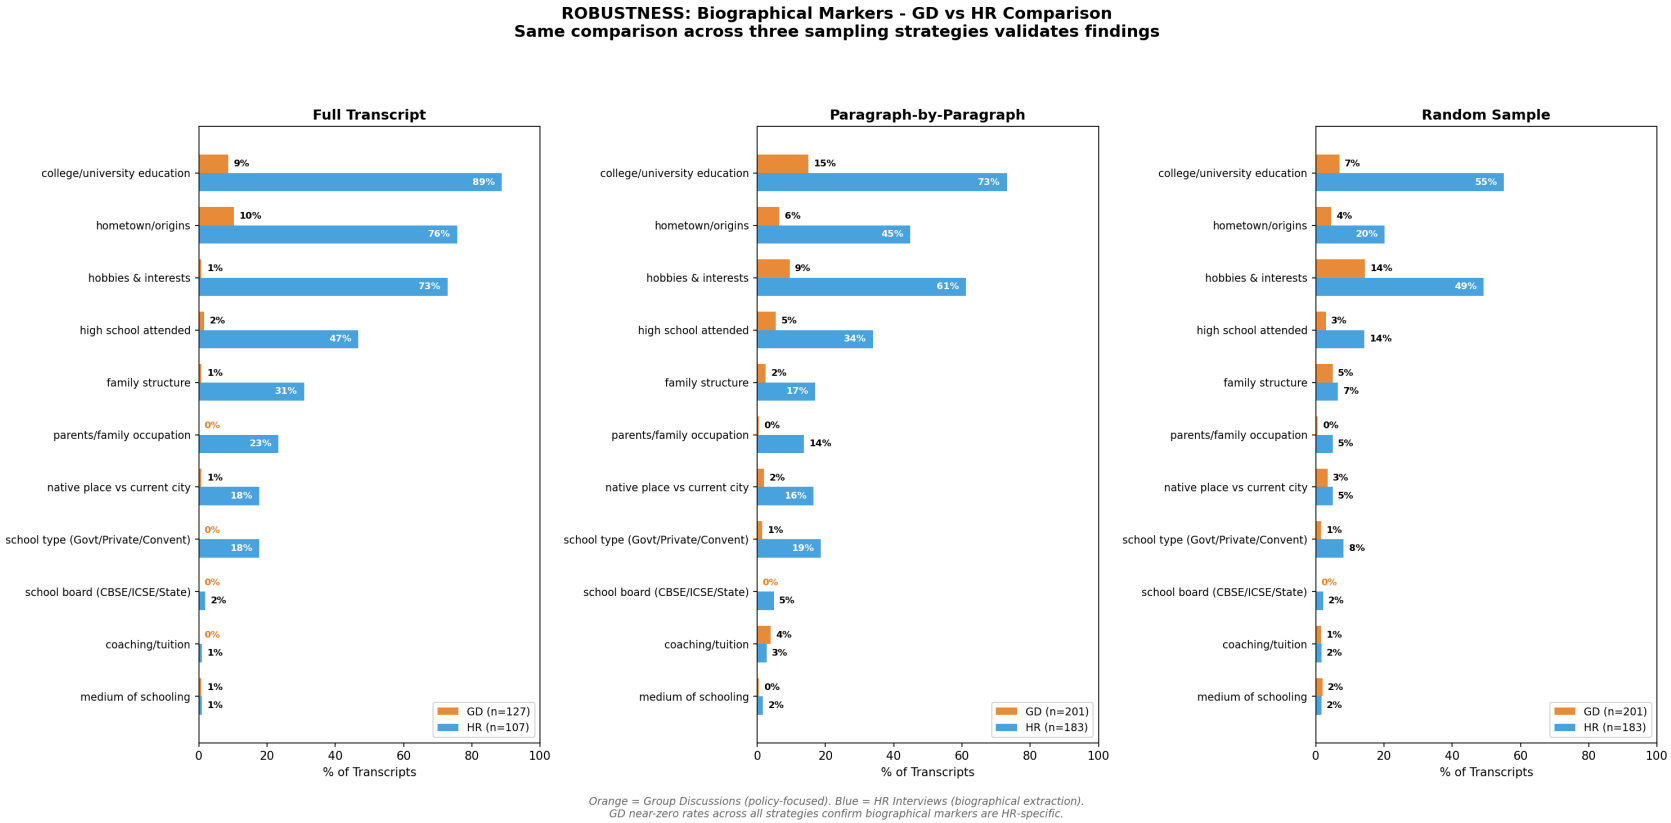


Figure J.7: **Claude Sonnet 4: Robustness Biographical Markers Across Sampling Strategies**

Notes: Figure J.7 reports the biographical-marker robustness check for Claude Sonnet 4 across full, paragraph, and random sampling strategies. Following the Appendix J convention, marker summaries exclude college-education prompts and focus on direct biographical extraction (for example, hometown, hobbies, family structure, and parents' occupation). The figure is estimated on Claude's own "live" subset; cross-model benchmark percentages appear in Table J.4.

GPT-5.2 Robustness Visuals

Figures J.8, J.9, and J.10 provide the same three robustness diagnostics for `gpt-5.2`. As in the Claude subsection, the goal is to show pattern stability within model across sampling strategies, using GPT-5.2's own "live" subset after model-specific filtering.

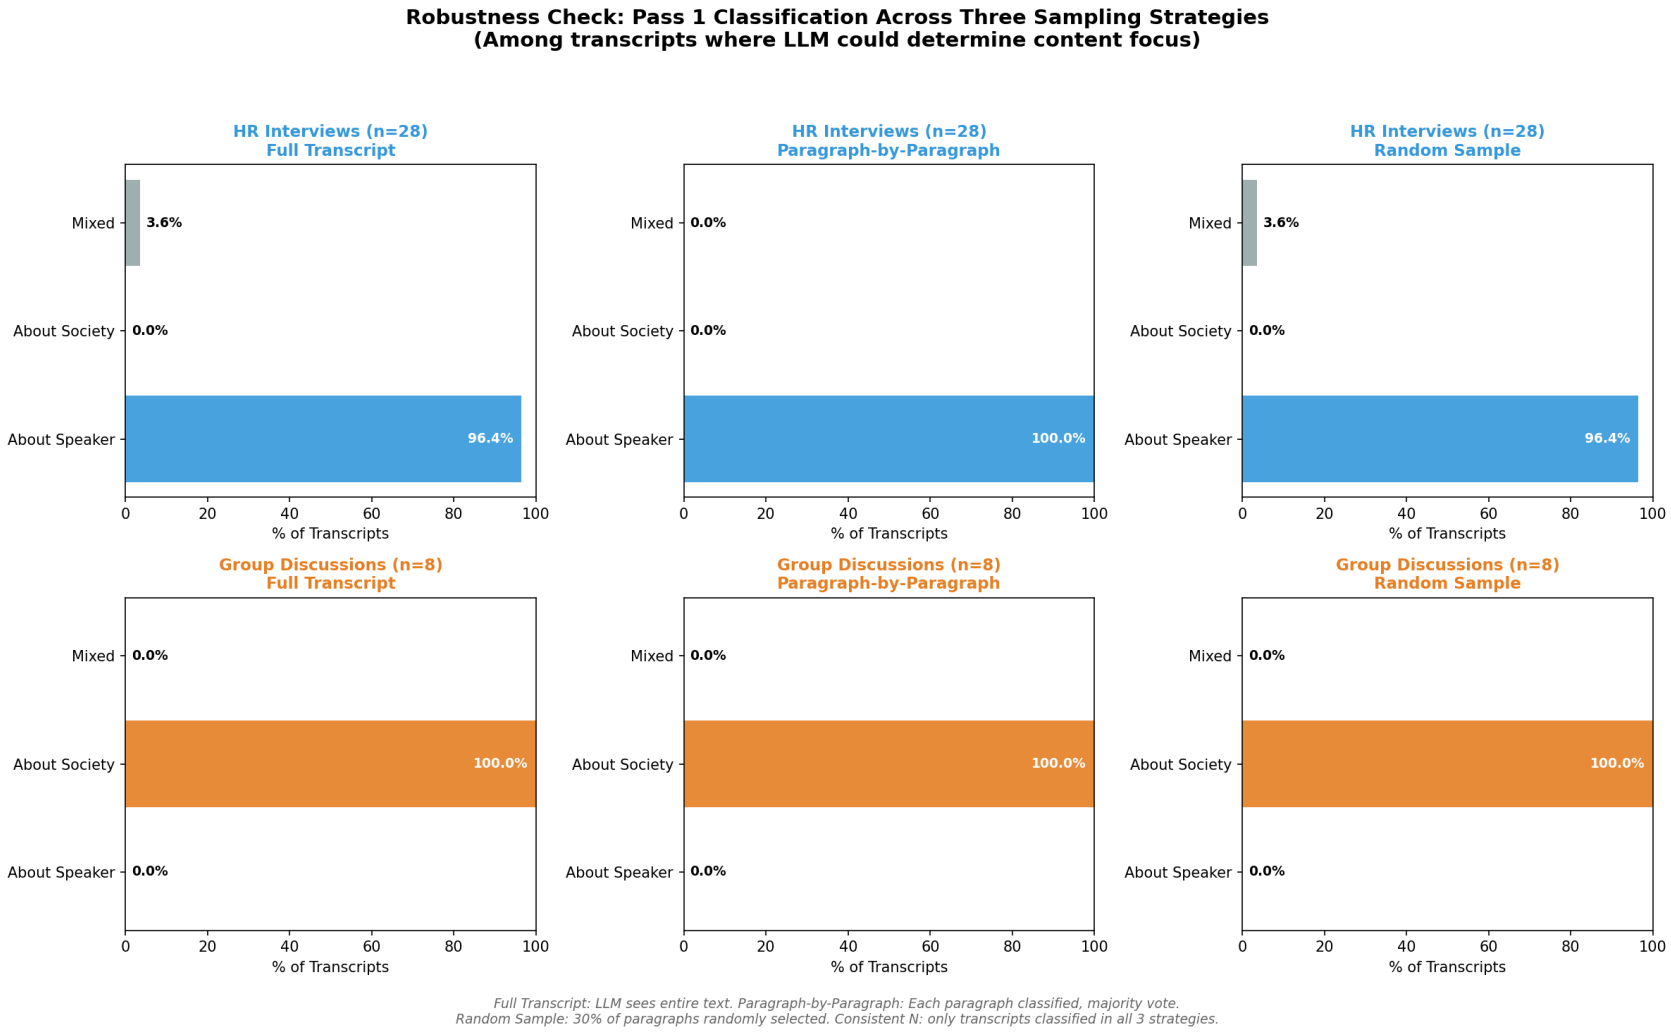


Figure J.8: **GPT-5.2: Robustness Content Focus Across Sampling Strategies**

Notes: Figure J.8 reports the content-focus robustness check for `gpt-5.2` under full-transcript, paragraph-by-paragraph, and random sampling strategies. Bars report transcript shares classified as about speaker, about society, or mixed, separately for HR interviews and group debates. The figure uses GPT-5.2's model-specific "live" subset and strategy-consistent classified transcripts. Table J.4 provides the corresponding cross-model percentage benchmarks.

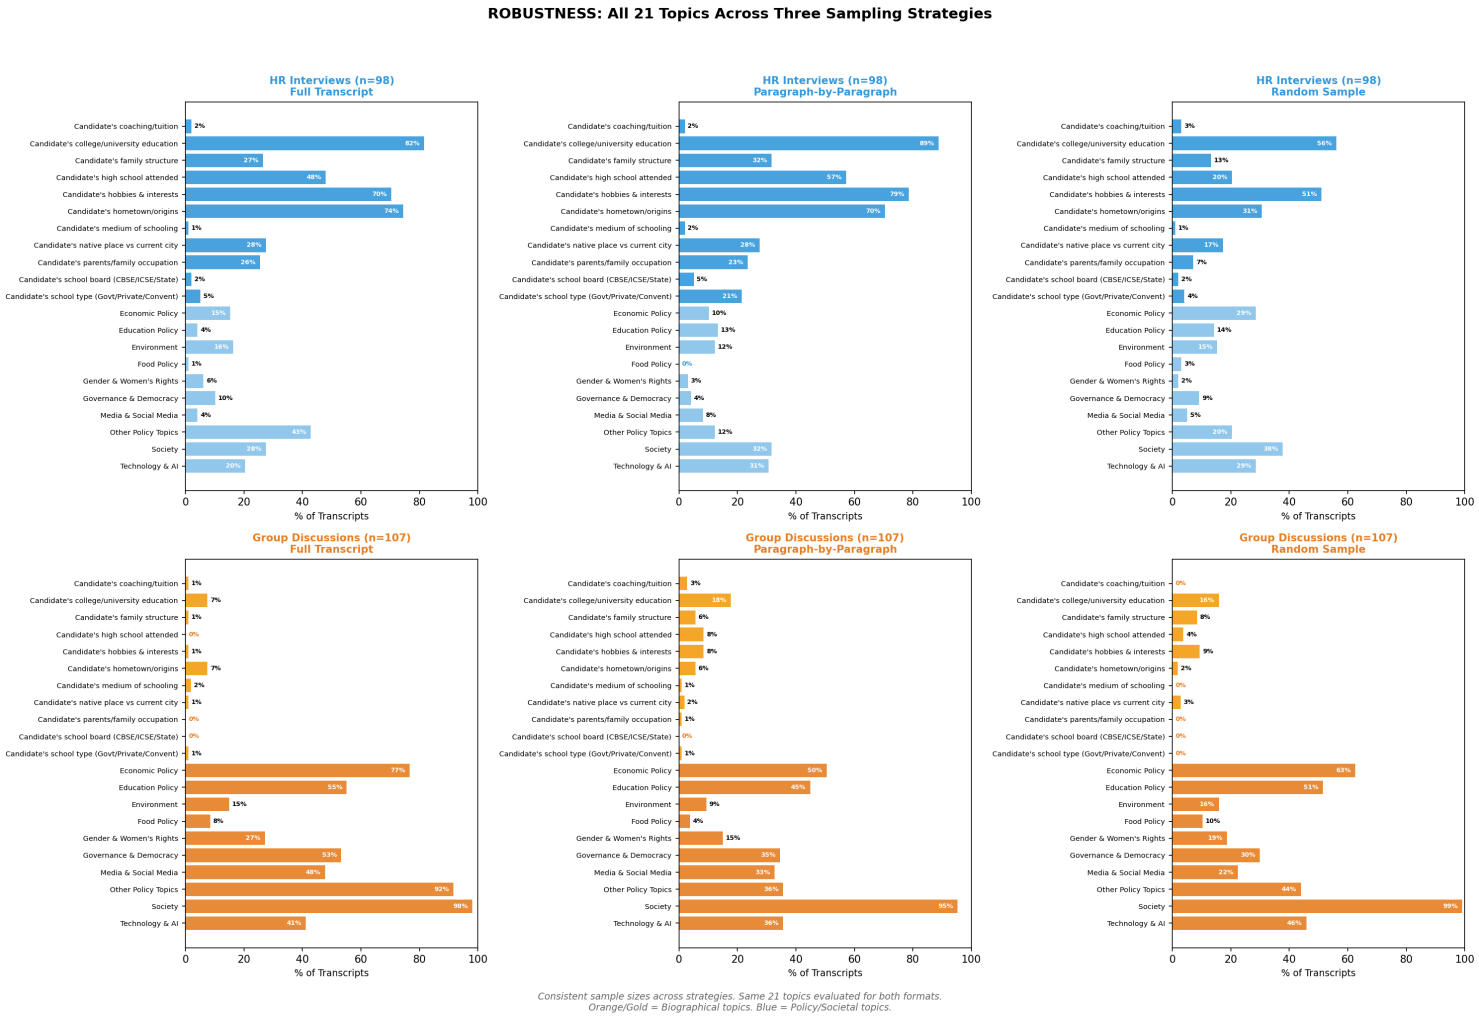


Figure J.9: **GPT-5.2: Robustness Topic Dimensions Across Sampling Strategies**

Notes: Figure J.9 reports the 20-topic robustness visualization for `gpt-5.2` (10 policy and 10 biographical topics) across full, paragraph, and random sampling strategies. Topic coding follows the unified taxonomy and prompt family documented in Appendix J. This visual check is designed to verify within-model topic-family separation across sampling strategies; Table J.4 summarizes model-to-model percentages.

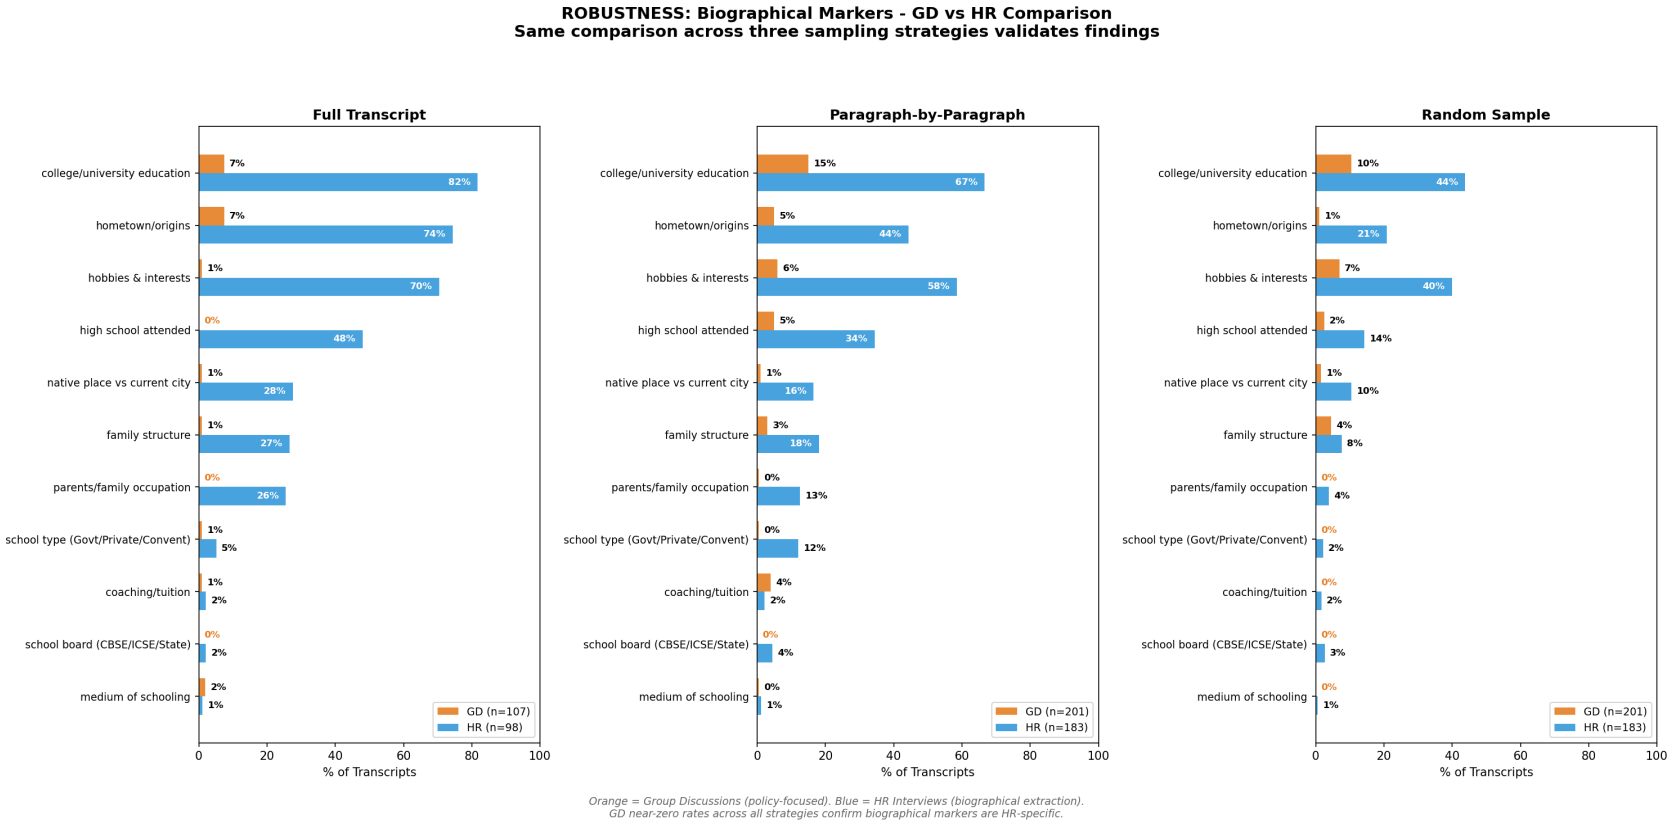


Figure J.10: **GPT-5.2: Robustness Biographical Markers Across Sampling Strategies**

Notes: Figure J.10 reports the biographical-marker robustness check for `gpt-5.2` across full, paragraph, and random sampling strategies. Consistent with the Appendix J definition, the marker summary excludes college-education prompts and captures only direct biographical probing. The figure is estimated on GPT-5.2's model-specific "live" subset; cross-model comparison values are reported in Table J.4.

Taken together, the Claude Sonnet 4 and GPT-5.2 visual replications preserve the same qualitative HR-versus-GD asymmetry as the baseline GPT-4o-mini run.

Validation Summary

The consistency across sampling strategies and across independent LLMs confirms that the content differences between HR interviews and group debates are fundamental and not sensitive to methodological choices. The full transcript strategy with GPT-4o-mini provides our primary estimates; paragraph-by-paragraph sampling, random sampling, and Claude/GPT-5.2 replications serve as robustness checks. Under all approaches, the core finding persists: HR interviews systematically probe biographical content while group debates focus on policy debates.

K Detailed Methodology for Classifying Caste-Identifiable Surnames

This appendix details the sources and procedures used to code student surnames as *caste-identifiable*. We define a surname as identifiable only when it unambiguously maps a candidate to an *administratively notified* community that the Indian state recognizes for reservation benefits. Our methodology proceeds in order of legal precision, prioritizing official government lists.

K.1 Hierarchy of Data Sources

We rely on the following sources, prioritized by their legal authority:

1. **Official central/state lists that *define* eligibility.** We begin with the authoritative government lists that legally define disadvantaged status:

    - ***OBC (Other Backward Classes):*** We use the National Commission for Backward Classes' state-wise *Central List of OBCs* (NCBC portal). For each State/Union Territory, clicking through reveals a table with a *Caste/Community* column and the corresponding *Gazette Resolution* numbers (e.g., Uttar Pradesh list).[19]
    - ***SC/ST (Scheduled Castes and Scheduled Tribes):*** Scheduled Castes (SC) and Scheduled Tribes (ST) are defined by Presidential Orders issued under Articles 341 and 342 of the Constitution and amended by Parliament. We use the Government of India's consolidated *List of Scheduled Castes* with the underlying constitutional orders (Ministry of Social Justice & Empowerment) and the *Constitution (Scheduled Tribes) Order, 1950* (Ministry of Tribal Affairs) along with the updated state-wise list (MoTA updated list, 2022).[20]

2. **State gazette notifications and implementing circulars.** Where needed to harmonize nomenclature (synonyms, regional titles, subgroup labels) with the central lists, we consult the underlying State/Centre *Gazette Resolutions* linked from the NCBC pages (see example above). These notices are the primary documents that effect inclusion and specify scope (e.g., area restrictions, sub-groups).

3. **Ethnographic cross-walks for name variants (used only to *corroborate* obvious cases).** To avoid false positives, we do not assign category based on ethnography alone. Instead, we use the Anthropological Survey of India's (AnSI) *People of India* series to cross-check widely used eponyms, aliases, and regional spellings for communities already recognized in (1)–(2) (AnSI homepage). This step verifies that a surname in our setting is a

[19] *Illustrative mapping (UP)*: The NCBC *Central List* for Uttar Pradesh lists *"Ahir, Yadav"* as OBC, with Gazette resolutions `12011/88/98-BCC` (06/12/1999) and `12011/68/93-BCC(C)` (10/09/1993). See the NCBC Uttar Pradesh page. "OBC" is the official term "Other Backward *Classes*," used in central policy; it overlaps with "Socially and Educationally Backward Classes (SEBC)" in constitutional usage. The NCBC *Central List* is the controlling reference for central-government OBC recognition.

[20] The Legislative Department also hosts the constitutional orders: *Constitution (Scheduled Castes) Order, 1950*: link.

commonly used label for the notified community, but it never substitutes for the legal lists.

K.2 Administrative Framework and the Definition of Disadvantage

What does "Caste/Community" on the NCBC pages denote? On the NCBC *Central List* pages, the *Caste/Community* column enumerates the officially notified communities eligible for OBC benefits in that State/Union Territory. In practice, these entries are typically endogamous birth-groups (*jātis*) or clusters of closely related *jātis*, sometimes listed with well-known aliases or titles. Our coding treats a surname as *identifiable* only if it is a widely recognized eponym for a notified community in the relevant State/Union Territory and that mapping is unambiguous across sources (1)–(2). Surnames shared across multiple communities/states or that do not provide a clear advantaged-vs-disadvantaged signal are coded *non-identifiable*.

How do community lists map to advantaged versus disadvantaged status? One might argue that these lists enumerate communities or jātis, not advantaged/disadvantaged status per se. This misunderstands India's administrative architecture. The very categorization into SC, ST, or OBC constitutionally defines disadvantage—these are groups the state recognizes as requiring affirmative action due to historical discrimination and socioeconomic marginalization. Those *not* appearing in any of these lists comprise the "General" or "Unreserved" category, which we code as advantaged. This binary—reserved (disadvantaged) versus unreserved (advantaged)—structures India's entire affirmative action system.

Consider concrete examples: "Chamar" appears in the SC list for multiple states, reflecting this leather-working community's historical untouchability and continuing disadvantage. Similarly, "Teli" (oil-pressers) appears in OBC lists across North India, recognized for educational and economic backwardness. Conversely, surnames like "Sharma" (Brahmins), "Agarwal" (merchant communities), or "Thakur" (Rajputs) appear in none of these lists, marking advantaged status. The state's administrative categorization thus provides the binary classification we employ: presence in SC/ST/OBC lists means they are classified as "disadvantaged"; absence means they are classified as "advantaged."[21]

K.3 Operational Mapping and Conservative Criteria

Mapping surnames to notified communities Operationally, we follow a two-step mapping: first link each surname to widely recognized community labels (using state/central alias lists, with AnSI only for corroboration), and then verify that those communities are notified as SC, ST, or OBC in the relevant state/central lists; only if this two-step mapping is unique and category-consistent do we flag a surname as *identifiable*. Critically, the government lists enumerate *communities*, not exhaustive surname inventories. A surname becomes identifiable when we can establish an unambiguous chain: Surname $\rightarrow$ Community $\rightarrow$ Administrative Category.

[21]The terms "forward castes," "upper castes," or "General category" all refer to communities not included in the constitutionally recognized disadvantaged categories. This binary—reserved versus unreserved—structures India's entire affirmative action system in education and public employment.

We follow a conservative rule: a surname is identifiable only if *both* (a) it is pan-regional (consistently maps to the same category across states), *and* (b) it is category-consistent—either mapping uniquely to one notified community, or to multiple communities all within the same disadvantaged category (SC/ST/OBC) and never to the General category. Any cross-category or state-dependent mapping is coded *non-identifiable*.

Appendix K, Table K.1 provides illustrative mappings under this conservative rule.

Table K.1: **Illustrative Surname-to-Category Mapping Under the Conservative Rule**

Surname (Illustrative)	Administrative Link Used for Coding	Coded as Identifiable?	Reason
Yadav	Linked to communities notified as OBC in central/state lists (e.g., NCBC state entries)	Yes	Pan-regional and category-consistent disadvantaged signal.
Chamar/Jatav	Linked to communities notified as SC under constitutional schedules	Yes	Unambiguous notified disadvantaged category mapping.
Patel	Used by multiple communities across states with non-unique category mapping	No	State-dependent and cross-community ambiguity.
Verma	Title used across multiple communities and categories in different regions	No	Cross-category ambiguity prevents a unique administrative mapping.

Notes: Table K.1 shows illustrative examples only. Coding in the data follows the full two-step procedure described above (surname → community label → notified administrative category) and uses official lists as the binding source. A surname is coded as identifiable only when mapping is pan-regional and category-consistent.

Accounting for Regional Variation Without Home State Information For instance, while "Kurmi" appears as an OBC community across North India, members often use surnames like "Patel" (in Gujarat), "Mahto" (in Bihar), or "Verma" (in UP). Similarly, the SC category includes "Chamar/Jatav," but members may use surnames like "Raigar," "Raidas," or "Ahirwar" depending on region and subgroup. However, since we do not observe students' home states or regions of origin, we code surnames as identifiable only when: (1) ethnographic and administrative sources confirm the surname is predominantly used by a specific notified community *across multiple states*, and (2) the mapping is unambiguous—the surname doesn't span multiple communities with different reservation statuses. This conservative approach—*requiring pan-regional clarity without state-specific information*—explains why only 15% of our sample has identifiable surnames. We exclude cases where surnames are shared across General (advantaged) and Non-General (disadvantaged) categories, where regional variation creates ambiguity, or where a surname's community affiliation is state-dependent.

Salience over completeness. Our approach prioritizes the surnames most salient to recruiters—those unambiguously linked to major notified communities that comprise substantial population shares. The government lists we use enumerate the largest disadvantaged

communities in each state, collectively accounting for SC (∼17%), ST (∼9%), and OBC (∼41%) of India's population. Surnames that clearly signal membership in these major communities—like "Yadav" for OBCs, "Chamar/Jatav" for SCs, or "Meena" for STs—are precisely those most likely to be "top of mind" for recruiters making rapid assessments. Critically, we map surnames to the binary advantaged/disadvantaged distinction relevant for hiring (presence/absence in SC/ST/OBC lists), not to specific jātis as ethnographic studies typically do. The People of India sources are used only to verify that a surname is indeed a recognized identifier for communities already in the government lists, never as a standalone basis for classification. While other caste-identifying surnames surely exist beyond our sources, our conservative approach captures the most salient cases where caste identification would be clearest. That only 15% of students have such unambiguously identifiable surnames—despite focusing on the most recognizable markers—underscores how limited surname-based caste identification is in practice, particularly without knowing students' home states.

L Differential Selection Across Surname Samples

A potential concern with the surname identifiability analysis is that the identifiable and non-identifiable samples may differ in ways that could explain the timing differences in caste penalties. In particular, if disadvantaged-caste candidates with identifiable surnames were systematically less qualified than their advantaged-caste peers (compared to the non-identifiable sample), this composition difference could generate an apparent early penalty that reflects skill gaps rather than discrimination.

Appendix L, Table L.1 addresses this concern at the mean level: the interaction between disadvantaged caste and non-identifiable surnames in the GPA regression is near zero, indicating no differential caste-related GPA slope across surname samples. Appendix L, Figures L.1 and L.2 then extend this check to the full distribution of job selectivity and show that the timing split persists after controls.

Table L.1: **Differential Selection Test: GPA by Caste and Surname Identifiability**

	log(GPA)
Disadv. Caste	-0.115***
	(0.010)
Non-Identifiable	-0.004
	(0.006)
Disadv. Caste × Non-Identifiable	0.001
	(0.011)
Observations	4,005

Notes: Table L.1 reports OLS estimates of log GPA regressed on caste status, surname identifiability, and their interaction. The sample includes all students who applied to at least one job (n=4,005). "Disadv. Caste" indicates Scheduled Caste, Scheduled Tribe, or Other Backward Classes status; "Non-Identifiable" indicates students whose surnames do not reveal caste affiliation under the conservative mapping rule. The near-zero interaction coefficient (0.001) indicates that the caste-GPA relationship is effectively the same across surname-identifiability groups, which weakens selection-based explanations for earlier penalties in the identifiable sample. Controls include department, degree, and cohort-year fixed effects; heteroskedasticity-robust standard errors are in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Figures L.1 and L.2 address this concern by comparing success rate gaps *without* and *with* controls for productivity-relevant traits. The key insight is that any differential selection on observable characteristics should show up as a difference between raw gaps and conditional gaps.

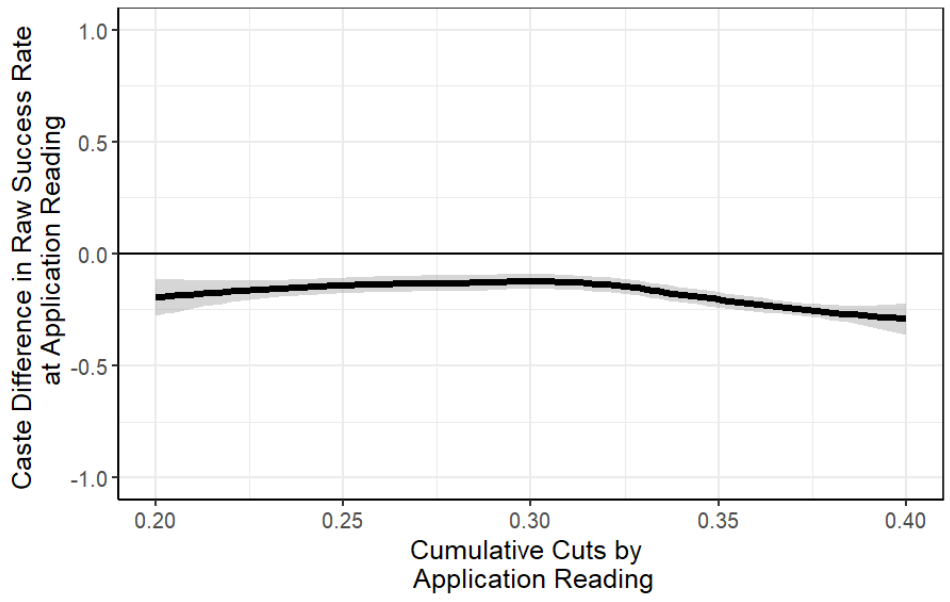
Figure L.1 shows *raw* success rate gaps—without controlling for GPA, entrance exam scores, or other observable traits. Both samples show substantial negative gaps, around -0.2 to -0.3 across job selectivity levels. This reflects the well-documented fact that disadvantaged-caste candidates have, on average, lower GPAs and test scores than advantaged-caste candidates in the same institution.

Figure L.2 shows *conditional* success rate gaps—after controlling for GPA, entrance exam score, 10th and 12th grade marks, department, degree, and cohort-year fixed effects. Here, the two samples diverge sharply:

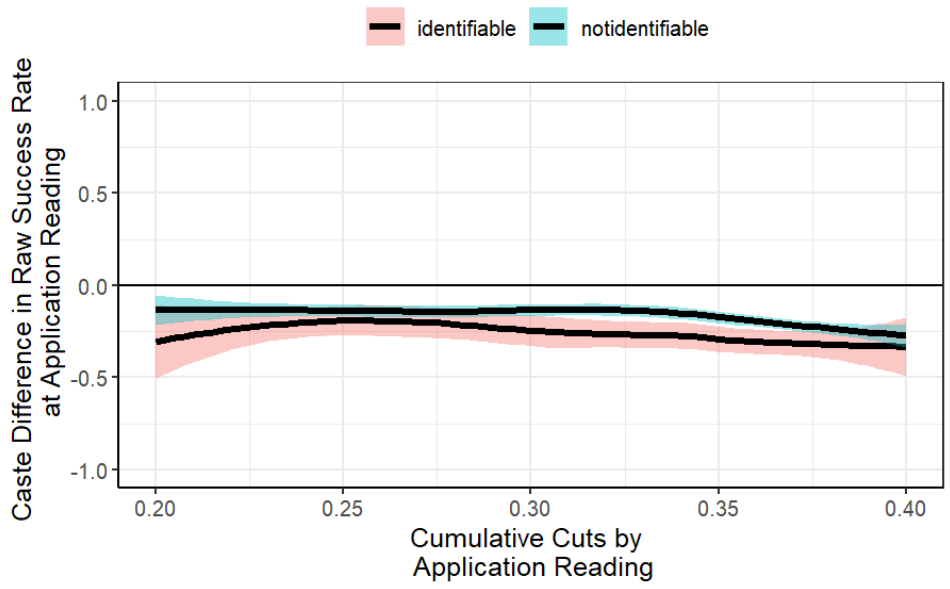
- **Non-identifiable surnames (cyan):** The gap rises to approximately zero across all job selectivity levels, with confidence intervals consistently crossing zero.
- **Identifiable surnames (red):** The gap remains persistently negative, around -0.1 to -0.15, even after conditioning on the full set of productivity-relevant controls.

This pattern has two important implications. First, for the non-identifiable sample, controlling for observable productivity traits fully accounts for the raw caste gap at the application stage. This means there is no differential selection on *unobservable* productivity traits that would generate a residual penalty—if there were, we would see a negative conditional gap even after controlling for observables. Second, for the identifiable sample, the negative gap *persists* after conditioning on observables, indicating that the penalty reflects something beyond differential productivity. Combined with the timing evidence—penalties emerging before any in-person interaction—the most parsimonious interpretation is that employers are responding to caste itself, as signaled by the surname.

Figure L.1: **Raw Caste Gap in Success Rates by Surname Identifiability (Without Controls)**



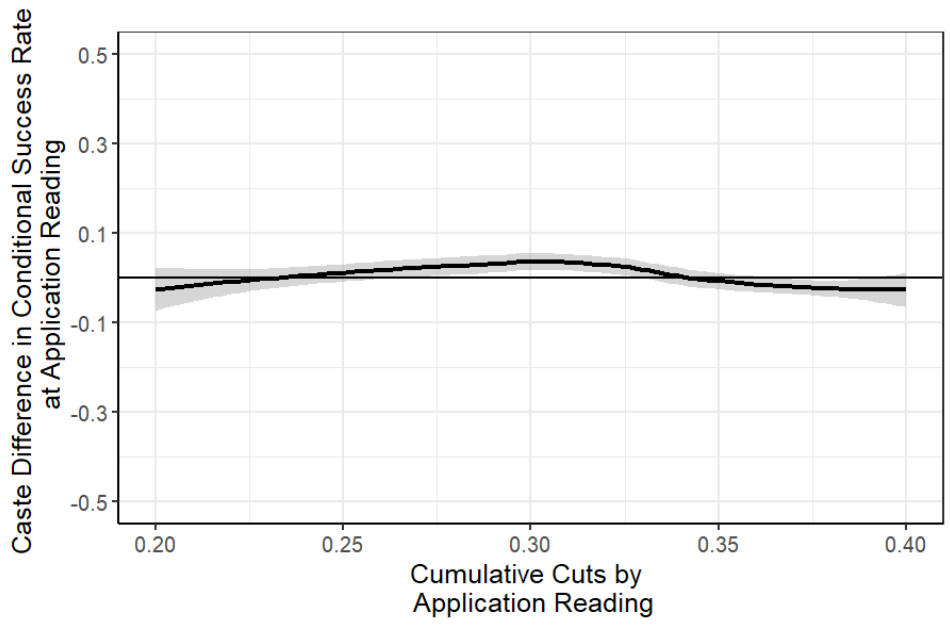
(a) Full Sample: Raw Caste Gap in Success Rates versus Cumulative Cuts



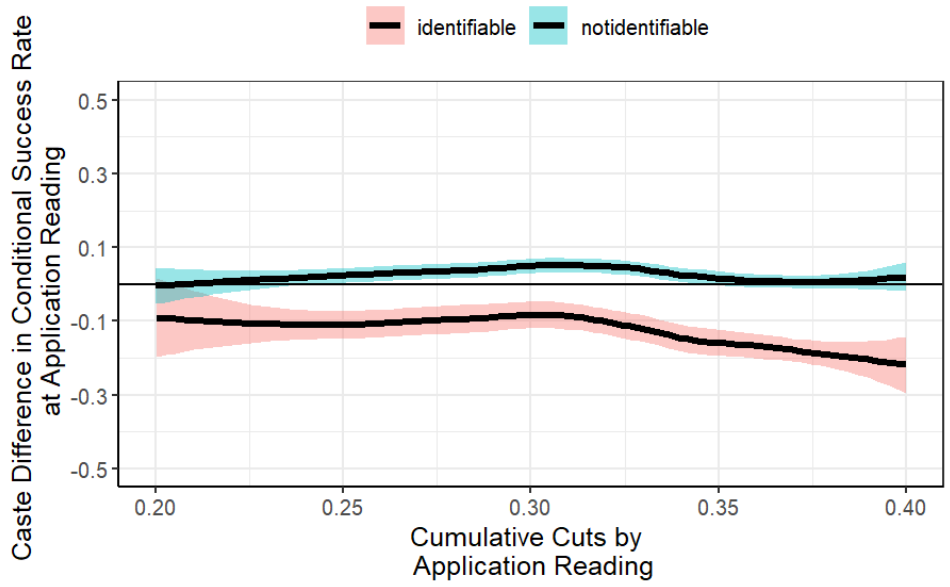
(b) Identifiable vs. Non-Identifiable Samples: Raw Caste Gap in Success Rates versus Cumulative Cuts

Notes: Figure L.1 shows *raw* caste differences in success rates at the application reading round, *without* controlling for productivity-relevant traits. Panel (a) shows the full sample; Panel (b) separates students with identifiable (red, n=619) and non-identifiable (cyan, n=3,386) caste surnames. The x-axis shows the cumulative proportion of applicants cut by each job. Y-values represent the raw difference in success rates between caste groups. Shaded areas show 95% confidence intervals; the horizontal line at zero indicates no caste difference. Both samples show negative gaps around -0.2 to -0.3, reflecting composition differences: disadvantaged-caste candidates have lower average GPAs and test scores than advantaged-caste candidates. Compare to Figure L.2: once observable productivity traits are controlled, the non-identifiable sample gap rises to approximately zero, while the identifiable sample gap remains persistently negative.

Figure L.2: **Conditional Caste Gap in Success Rates by Surname Identifiability (With Controls)**



(a) Full Sample: Conditional Caste Gap in Success Rates versus Cumulative Cuts



(b) Identifiable vs. Non-Identifiable Samples: Conditional Caste Gap in Success Rates versus Cumulative Cuts

Notes: Figure L.2 shows *conditional* caste differences in success rates at the application reading round, after controlling for GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, department, degree, and cohort-year fixed effects. Panel (a) shows the full sample; Panel (b) separates students with identifiable (red, n=619) and non-identifiable (cyan, n=3,386) caste surnames. The x-axis shows the cumulative proportion of applicants cut by each job. Y-values represent the disadvantaged caste coefficient from job-specific linear probability models with the full set of controls. Smoothed conditional means shown with 95% confidence intervals (shaded); the horizontal line at zero indicates no caste difference. Compare to Figure L.1: once observable productivity traits are controlled, the non-identifiable sample gap rises to approximately zero, while the identifiable sample gap remains persistently negative. This indicates no differential selection on unobservable productivity traits for the non-identifiable sample.

M More Details on Alternative Explanations

Several alternative explanations are plausible *a priori*: cumulative ability deficits that emerge only after early screens, statistical discrimination based on noisy signals, customer discrimination in client-facing roles, strategic delay of discrimination until late stages, or selection on unobservables correlated with caste.

Taken together, these alternatives are difficult to reconcile with the timing and post-hire patterns observed. In the pooled sample, pre-interview stages show near-zero gaps even in highly selective jobs; if interviews revealed true productivity deficits, the promotion advantage for disadvantaged hires should shrink rather than expand. Gaps appear in both client-facing and non-client-facing roles, undermining customer-discrimination accounts. Surname-identifiability patterns indicate that penalties track information revelation rather than gradual accumulation. Oster-style bounds and alternative functional forms further limit omitted-variable interpretations (Oster, 2019).

1. **Early-round "fairness."** If firms were lowering early-stage cutoffs to appear inclusive, disadvantaged candidates should display *higher* advancement rates in early rounds, especially in highly selective jobs. Instead, pre-interview gaps are near zero in the pooled sample, and candidates with caste-identifiable surnames face penalties already at application reading in the success-rate split (Figure L.2).

2. **Delayed discrimination for optics.** Firms have no obvious incentive to postpone discrimination in a setting without caste-focused regulatory oversight; when caste is visible through surnames, penalties appear immediately. If bias were strategically delayed, group debates—already subjective—would be a natural stage to implement it, yet no gaps appear there. Qualitative work also documents firms' general resistance to caste-based affirmative action (Jodhka, 2008; Jodhka and Newman, 2007).

3. **Hidden right-tail skill gaps.** If disadvantaged candidates had weaker right-tail ability, the gap should emerge in highly selective jobs *before* interviews. This pattern does not appear: pre-interview gaps are small even in the most selective jobs, while post-hire performance advantages persist.

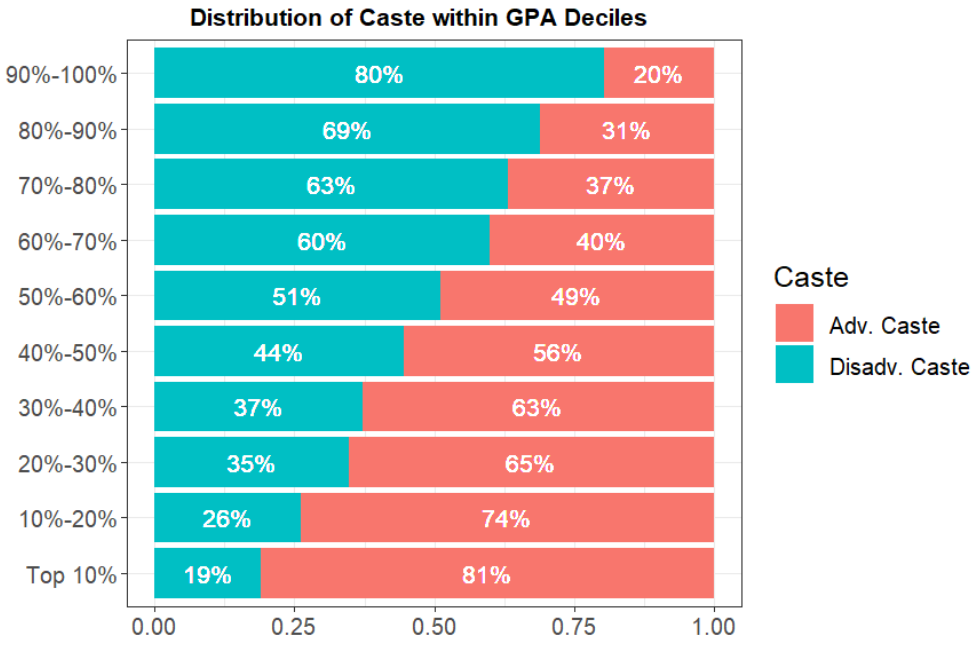
4. **Socio-emotional skill differences.** Group debates explicitly assess communication and teamwork. The absence of caste gaps in debate outcomes is inconsistent with socio-emotional skill deficits that would explain interview-stage penalties.

5. **Outside options or missing data.** The placement office enforces mandatory participation once students register, prohibits skipping rounds, and records every exit decision through the centralized portal (Sections 4 and 6.1). Exit surveys corroborate the records, limiting scope for unobserved attrition.

6. **Preferential promotion for disadvantaged hires.** Caste-focused diversity initiatives in the private sector are rare: only a small fraction of major firms track caste explicitly, and qualitative evidence documents resistance to caste-based affirmative action (Bapuji et al., 2023; BusinessLine, 2018). The DEI analysis in Appendix C shows caste is almost absent from formal policy, making preferential promotion an unlikely driver of the observed advantage.

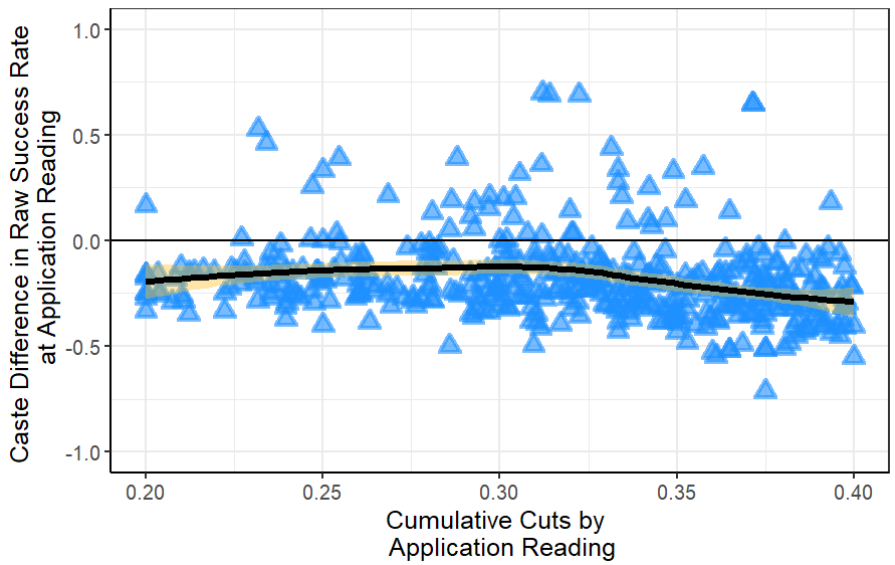
N Additional Figures

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

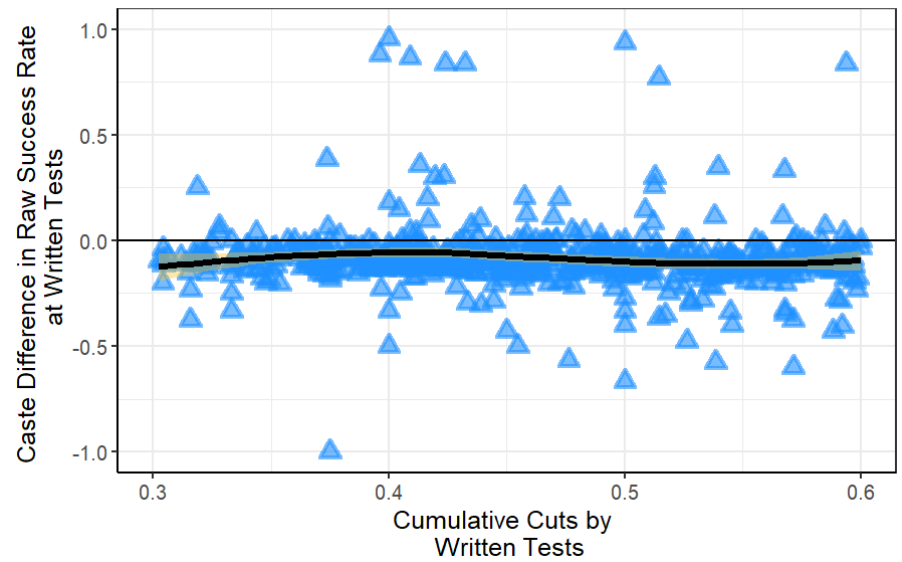


Notes: Figure N.1 demonstrates "common support"—the presence of both advantaged and disadvantaged caste students at every level of academic performance. Each bar shows the count of students within a GPA decile, with advantaged caste students in one color and disadvantaged caste students in another. GPAs are normalized to zero mean and unit standard deviation within each degree-year. The overlap across all deciles ensures that comparisons between caste groups are not confounded by differences in the support of the GPA distribution—i.e., both groups are represented throughout the ability range, including the right tail. This overlap is important for the paper's identification strategy, which relies on comparing outcomes across jobs with varying selectivity thresholds.

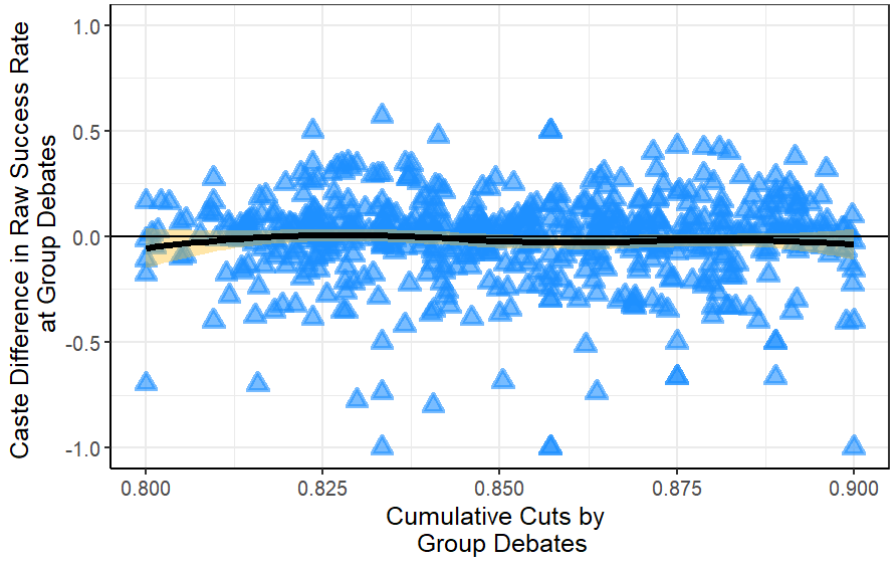
Figure N.2: **Caste Penalty in the Uncontrolled Success Rate At Each Screening Round**



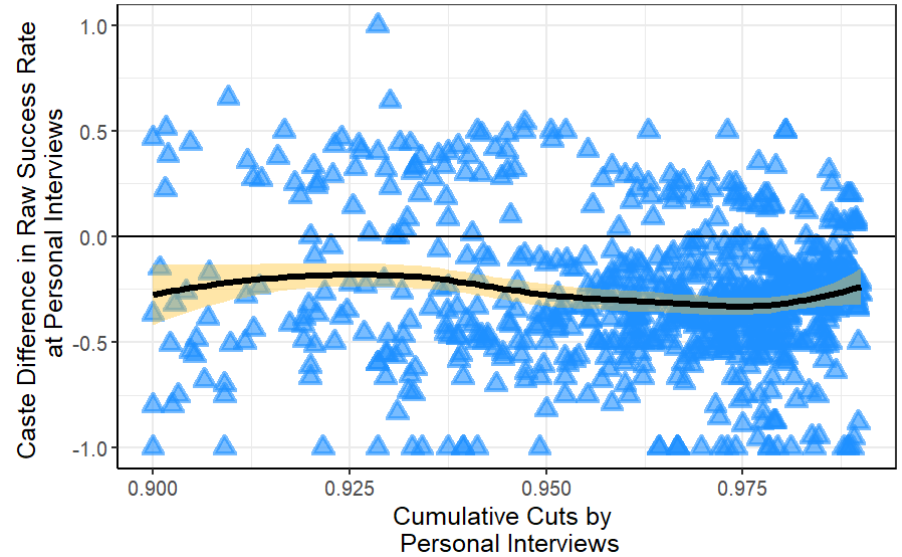
(a) Caste Penalty in Raw Success Rates versus Cumulative Cuts by Application Reading



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



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



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

Notes: Figure N.2 plots *raw* (uncontrolled) caste differences in success rates against job selectivity at each of four screening stages: application reading, written tests, group debates, and personal interviews. Each blue triangle represents a single job. Unlike Figure 3, these are simple differences in mean advancement rates between disadvantaged castes (SC/ST/OBC) and advantaged castes—with no individual-level controls and no fixed effects. The y-axis shows the raw caste gap for each job; negative values indicate lower advancement rates for disadvantaged castes. The x-axis shows job selectivity, measured as the cumulative proportion of applicants eliminated by that stage. Black lines show smoothed conditional means (LOESS); shaded regions show 95% confidence intervals. The horizontal line at zero indicates no caste difference. Key finding: a modest raw caste gap exists at the application stage, but this gap *shrinks toward zero* as selectivity increases through written tests and group debates. The gap then *re-emerges sharply* at interviews. This shrinking-then-reappearing pattern—observed even without controls—is inconsistent with cumulative ability deficits and suggests positive selection of disadvantaged candidates reaching interviews.

O Additional Tables

Table O.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 O.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.

Table O.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 O.2 reports pre-college skills (entrance exam scores, 10th and 12th standard national exit exam scores) and college GPA by caste and degree. All scores normalized to zero mean and unit standard deviation within each degree and year. Entrance exam scores (originally ranks) renormalized so higher numbers are better. The "Caste Diff." column shows the difference (disadvantaged minus advantaged) in standard deviation units, with statistical significance indicated by asterisks. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$. Adv. Caste and Disadv. Caste denote advantaged and disadvantaged castes.

Table O.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 O.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. The firms in the sample, regardless of headquarters location, are predominantly multinational corporations.

Table O.4: **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 O.4 reports mean and median candidate counts per job at each stage of the four-round screening process: applications received, candidates advancing to written tests, candidates advancing to group debates, candidates advancing to personal interviews, and final offers extended. Values calculated by averaging student counts across years per job, then averaging across jobs. The declining counts illustrate the substantial selectivity at each screening stage.

Table O.5: **Caste Differences in the Success Rate at Each Job Search Stage by Job Type**

Panel A. Client-Facing Jobs

	App → Test (1)	Test → Debate (2)	Debate → Int (3)	Int → Offer (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	✓	✓	✓	✓

Panel B. Non-Client-Facing Jobs

	App → Test (1)	Test → Debate (2)	Debate → Int (3)	Int → Offer (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: Table O.5 reports within-job linear probability models of advancing to the next hiring stage by job type. Panel A restricts to client-facing jobs; Panel B restricts to non-client-facing jobs. Coefficients are shown for both baseline OLS and Oster-adjusted specifications at each stage (application to test, test to debate, debate to interview, interview to offer). The dependent variable is an indicator for advancing, with observation at the student-job-year level conditional on remaining in contention. Controls include college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, and prior internship sector, with job fixed effects, department fixed effects, degree stream fixed effects, and cohort-year fixed effects. Bias-adjusted coefficients follow Oster (2019), setting $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$. Standard errors in parentheses. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

Table O.6: **Sample Selection into Success-Rate and Job-Quality Analyses: GPA Comparisons**

Panel A. Success-Rate Sample Selection (Participants vs. Non-Participants)

	Dependent variable: log (GPA)	
	(1)	(2)
Disadv. Caste	−0.114***	0.009***
	(0.004)	(0.002)
Did Not Participate in Job Fair	0.011	
	(0.016)	
Did Not Participate in Job Fair * Disadv. Caste	−0.044**	
	(0.022)	
Constant	2.071***	2.199***
	(0.010)	(0.028)
Observations	4,164	2,527
Adjusted R^2	0.231	0.822
Year FE	✓	✓
Major	✓	✓
Degree	✓	✓
Job FE		✓
Note:	*p<0.1; **p<0.05; ***p<0.01	

Panel B. Job-Quality Sample Selection (Placed vs. Not in Job-Quality Regressions)

	Dependent variable: log (GPA)	
	(1)	(2)
Disadv. Caste	−0.032***	0.009***
	(0.004)	(0.002)
Not in Job Quality Regression	−0.158***	
	(0.005)	
Not in Job Quality Regression * Disadv. Caste	−0.070***	
	(0.007)	
Constant	2.127***	2.199***
	(0.008)	(0.028)
Observations	4,164	2,527
Adjusted R^2	0.559	0.822
Year FE	✓	✓
Major	✓	✓
Degree	✓	✓
Job FE		✓
Note:	*p<0.1; **p<0.05; ***p<0.01	

Notes: Table O.6 reports GPA-based selection checks for the two core outcomes. Panel A examines selection into the success-rate analysis: Column (1) compares participants in the centralized job fair with non-participants, and Column (2) benchmarks within-job GPA among hires. Panel B applies the same structure for selection into the job-quality analysis by comparing students in versus out of job-quality regressions. Across panels, interaction terms show disadvantaged candidates excluded from the main regressions have lower GPAs, while disadvantaged hires within the same job have slightly higher GPAs. This pattern implies that exclusion is unlikely to mechanically inflate estimated caste penalties. Year, major, and degree fixed effects are included in both columns; job fixed effects are included in Column (2) only. *p<0.1; **p<0.05; ***p<0.01.

Table O.7: **Caste Differences in Job Quality 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: Table O.7 reports the coefficient on disadvantaged caste (SC/ST/OBC indicator) from OLS regressions of log(average salary) at each job search stage. Each column represents a different stage: (1) application, (2) written test, and subsequent rounds through job choice. 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 linear controls: college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, prior internship sector, department fixed effects, degree stream fixed effects, and cohort-year fixed effects. Job fixed effects are omitted because salaries are standardized and identical within jobs. Bias-adjusted coefficients follow Oster (2019), setting $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$. Standard errors in parentheses. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

Table O.8: **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: Table O.8 reports the coefficient on disadvantaged caste (SC/ST/OBC indicator) from OLS regressions of log(salary) of the accepted job, for graduates who obtained jobs through the college's job fair. Column (1) shows raw estimates without controls. Column (2) adds linear controls: college GPA, entrance exam score (JEE rank percentile), 10th and 12th grade marks, prior internship sector, department fixed effects, degree stream fixed effects, and cohort-year fixed effects. Columns (3)-(4) show separate estimates for client-facing and non-client-facing jobs with full controls. Job fixed effects are omitted because salaries are standardized and identical within jobs regardless of candidate characteristics. Bias-adjusted coefficients follow Oster (2019), setting $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$. Standard errors in parentheses. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

Table O.9: **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 O.9 shows the distribution of job selectivity across the four-round screening process. *Conditional cuts* report the fraction of candidates eliminated at each stage, given that they reached that stage. *Cumulative cuts* report the total fraction of original applicants eliminated by that stage. For example, by the personal interview stage, the median job has cumulatively eliminated 97% of original applicants (Q2 = 0.97 for Cumulative Cut at Personal Interview). Q1, Q2 (median), and Q3 denote quartiles across all jobs in the sample. This variation in cumulative selectivity is exploited in Section 6.2 to test whether caste penalties vary with job selectivity.

Table O.10: **Caste Differences in GPA Between Those Who Left the First Job Versus Those Who Did Not**

	Dependent variable:
	log (GPA)
Disadv. Caste	0.007^{***}
	(0.002)
Left Job	-0.003^{*}
	(0.002)
Left Job × Disadv. Caste	-0.004
	(0.003)
Observations	2,489
Adjusted R^2	0.832
Year FE	✓
Major	✓
Degree	✓
Job FE	✓
Note:	*p<0.1; **p<0.05; ***p<0.01

Notes: Table O.10 tests for differential selection on GPA between job leavers and stayers by caste. The dependent variable is log(GPA). The coefficient on Disadv. Caste indicates that disadvantaged hires have higher GPAs than advantaged peers within the same job. The coefficient on Left Job indicates that among advantaged workers, those who leave have slightly lower GPAs than stayers. The interaction term (Left Job × Disadv. Caste) tests whether the GPA gap between leavers and stayers differs by caste—the near-zero and insignificant coefficient indicates no differential selection. The regression includes job fixed effects, year fixed effects, and controls for major and degree. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table O.11: **Caste Differences in Job Quality at Each Job Search Stage by Surname Identifiability**

Panel A. 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						

Panel B. 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: Table O.11 reports stage-specific job-quality regressions by surname-identifiability sample. Panel A uses students with identifiable caste surnames; Panel B uses students with non-identifiable surnames. In both panels, columns track the job-search pipeline from application through job choice, and the dependent variable is log average posted salary across jobs still in contention (or accepted salary in the "Choice" column). All regressions include the main controls and cohort-year fixed effects; job fixed effects are omitted because salary is standardized within jobs. Bias-adjusted coefficients follow Oster (2019), with $\delta = 1$ and $R^* = 1.3 \times \tilde{R}$. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table O.12: **Caste Differences in the Success Rate at Each Job Search Stage by Surname Identifiability**

Panel A. Identifiable Sample

	App → Test	Test → Debate	Debate → Int	Int → Offer
	(1)	(2)	(3)	(4)
Disadv. Caste	-0.1839***	0.1099***	-0.0629***	-0.2328***
	(0.0037)	(0.0045)	(0.0071)	(0.0208)
Mean Success Rate	0.4640	0.8830	0.2080	0.4070
Observations	46,163	21,407	18,899	3,932
Adjusted R^2	0.62095	0.29942	0.04530	0.18969

Panel B. Non-Identifiable Sample

	App → Test	Test → Debate	Debate → Int	Int → Offer
	(1)	(2)	(3)	(4)
Disadv. Caste	0.0260***	-0.0112***	0.0078***	-0.2979***
	(0.0015)	(0.0020)	(0.0029)	(0.0050)
Mean Success Rate	0.6240	0.8010	0.2660	0.2450
Observations	259,378	161,743	129,590	34,446
Adjusted R^2	0.63019	0.35001	0.01739	0.24959

Notes: Table O.12 reports stage-specific linear probability models of advancing to the next hiring stage by surname-identifiability sample. Panel A uses students with identifiable caste surnames; Panel B uses students with non-identifiable surnames. Columns track the hiring pipeline from application reading through final interviews, and each coefficient reports the disadvantaged-caste gap in conditional progression at that stage. All specifications include the main linear controls, job fixed effects, department fixed effects, degree stream fixed effects, and cohort-year fixed effects. Standard errors are in parentheses. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

Supplementary References

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.

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

Harrell, F. E. (2015). *Regression Modeling Strategies: With Applications to Linear Models, Logistic and Ordinal Regression, and Survival Analysis* (2nd ed.). Cham, Switzerland: Springer.

Hobbs, A. (2014). *A Chosen Exile: A History of Racial Passing in American Life*. Cambridge, MA: Harvard University Press.

Ignatiev, N. (1995). *How the Irish Became White*. New York: Routledge.

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

Moorjani, P., K. Thangaraj, N. Patterson, M. Lipson, P.-R. Loh, P. Govindaraj, B. Berger, D. Reich, and L. Singh (2013). Genetic evidence for recent population mixture in india. *American Journal of Human Genetics 93*(3), 422–438.

Omi, M. and H. Winant (2014). *Racial Formation in the United States* (3rd ed.). New York: Routledge.

Schwartz, C. R. (2013). Trends and variation in assortative mating: Causes and consequences. *Annual Review of Sociology 39*, 451–470.

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*.

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.