Making the Elite: Coded Discrimination at Top Firms

Soumitra Shukla*

April 24, 2024

How quantitatively important are screening practices in driving socioeconomic gaps in access to elite jobs, and what aspects of screening generate such disparities? To answer this question, I use recruitment data from over 1000 jobs at elite multinational corporations, primarily based in the United States or Europe, that hire from an elite Indian college. Each employer conducts four eliminatory screening rounds. I show that nearly 90% of the caste disparities in hiring arise in the final screening round, which comprises informal, non-technical personal interviews that often assess "fit." Thus, the three screening rounds prior to personal interviews together explain only a modest portion of hiring disparities. These three rounds are application reading, written aptitude tests, and large group debates that assess socio-emotional skills. The caste penalty following personal interviews closely parallels caste revelation. In a small subsample where last names provide clearer caste signals, the penalty emerges much earlier—during the application reading round. However, the overall caste penalty during application reading is negligible because 85% of the sample has non-distinctive caste names. Instead, the caste penalty emerges during personal interviews, where students are screened based on background, hobbies, and cultural fit: characteristics strongly correlated with caste. Furthermore, caste disparities in hiring outcomes are not explained by differences in job productivity. Disadvantaged castes have about 15% higher promotion rates, even among the most selective jobs. By studying the entire entry-level hiring process among elite jobs, this paper shows that disparities in hiring predominantly arise in the final screening round—personal interviews—something that audit studies would not capture. My findings elucidate the economic behavior underlying elite hiring outcomes and pave the way for policies that reduce socioeconomic disparities in access to elite jobs by effectively targeting the source of the inequality, in India and elsewhere.

JEL codes: J70, J71, M50, M51

---

*Shukla: Economist, Board of Governors of the Federal Reserve System (soumitra.shukla@frb.gov) and Research Affiliate, Department of Economics, Yale University (soumitra.shukla@yale.edu). This paper was previously circulated under the title "Making the Elite: Top Jobs, Disparities, and Solutions." I would especially like to thank many members of the educational institution I worked with for immense support, motivation, and encouragement. 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 Alex Albright, Peter Arcidiacono, Crémieux, Alessandra Fenizia, Claudia Goldin, Peter Hull, Yujung Hwang, Patrick Kline, Kala Krishna, Jack Mountjoy, Eliza McNay, Dylan Moore, Tom Mroz, Cheyenne Quijano, Evan Rose, Kamalini Ramadas, Rob Valletta, Emmanuel Yimfor, Basit Zafar, Seth Zimmerman, seminar participants at the Yale labor workshop, the Chicago FED, the Board of Governors of the Federal Reserve System, Wharton (BEPP), Penn State, UCSD, SOLE, Southeastern Micro Labor Workshop, System Equitable Growth Conference at the New York Fed, and the "Discrimination in the 21st Century Conference" organized by the Becker Friedman Institute and Chicago Booth for helpful feedback and support. This paper was also selected to be presented at the Barcelona GSE Summer Forum 2022 (Structural Microeconometrics), which I was unable to attend due to COVID-related issues. I gratefully acknowledge financial support from the Cowles Foundation for Research in Economics and a Yale Dissertation Fellowship. The views in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Fed System.

1 Introduction

Screening practices commonly impact socioeconomic gaps in access to elite jobs and colleges (Stevens, 2009; Rivera, 2015). Such screening practices are often outwardly blind to race, gender, ethnicity, or social class. However, they can implicitly penalize certain groups by screening on cultural fit, personal hobbies, legacy and athlete preferences, and subjective impressions of personality (Rivera, 2012; Arcidiacono et al., 2021; SFFA v. Harvard, 2023; Chetty et al., 2023). Therefore, such screening practices can create barriers to elite jobs and colleges (henceforth, "elite attainment").

The pivotal socioeconomic role of elite attainment has led many governments, firms, and educational institutions around the world to embrace improving the economic mobility of historically disadvantaged groups as one of their main objectives (Hoxby and Avery, 2013; Zimmerman, 2019; Chetty et al., 2023). In India, most elite educational institutions have implemented government mandated, caste-based affirmative action policies (quotas) since the early 1990s, assigning 50% of their college seats to disadvantaged castes. Despite progress, substantial disparities persist. Stark caste disparities remain a feature of elite jobs: nearly 95% of board members in India's top 1,000 businesses belong to advantaged castes, which constitute only about 30% of the overall population (Dayanandan et al., 2019).

The lagging representation of disadvantaged castes in elite jobs, despite broad-based affirmative action in university admissions in India, could be partly driven by screening processes which allow employers to discriminate against disadvantaged castes. Even at American workplaces, screening practices have limited the economic mobility of disadvantaged castes, often by allowing managers to covertly retaliate against employees (California v. Cisco Systems, Inc., 2022). Yet, little is known about how quantitatively important screening processes are in driving caste disparities in access to the "elite," or what aspects of candidate screening drive these disparities. Adding to these challenges is the fact that recruitment practices of elite jobs are often non-transparent (Rivera, 2015). My paper addresses these issues by showing that the final screening round consisting of subjective, non-technical personal interviews that often assess "fit" could create barriers for disadvantaged castes in accessing elite jobs. Specifically, I show that nearly 90% of the caste disparities in hiring arise due to personal interviews. I also argue that this screening round is where employers are most likely to learn about caste. Further, information learned during personal interviews is weakly correlated with job performance. Instead, personal interviews, a common component of elite job recruitment in India and elsewhere, are a large contributor to unwarranted socioeconomic (caste) disparities in the labor market. Such disparities could be motivated by taste-based, statistical, customer, or coworker discrimination, or other factors. To my knowledge, this paper is the first to quantify *how much* of the hiring disparities, however motivated, among elite jobs in a real world and high-stakes setting, could be attributed to subjective assessments of "fit" done through personal interviews.

To uncover what aspects of screening drive the caste gap in hiring for elite jobs, I use recruitment data from elite multinational corporations (MNCs), predominantly based in the United States or Europe, that hire from an elite Indian college. Many of these firms regularly feature in the Fortune 500. Because nearly all elite Indian colleges conduct similar job fairs with participation from nearly the same group of firms, my paper effectively offers a representative window into how elite college graduates transition into "elite jobs," defined as formal entry-level jobs in MNCs, major national corporations, or unicorn firms hiring in the Indian private sector (Section 4.2).[1] Elite jobs comprise 98% of all jobs in my sample, with a majority

[1] A unicorn firm is typically a privately-owned startup with a valuation of over $1 billion.

headquartered in the United States or Europe. Almost all jobs in my sample hire for their Indian offices and pay in the top 1% of entry-level job salaries in the Indian private sector (Section 5.2).

Despite the elite college assigning nearly 50% of its seats to disadvantaged castes, there are large caste disparities in labor market outcomes. Caste is reported by students, typically by presenting a formal "caste certificate" to the college and is used as the basis for quota-based policies in admission. In this paper, "caste earnings gap" is defined as the average earnings of disadvantaged castes minus the average earnings of advantaged castes. The caste gap in earnings among graduates from the elite college is 23% without controls. In the presence of detailed controls for pre-college skills, college academic performance, previous labor market experience, and other employer-relevant skills, the gap is reduced to 10%. There are no differences in pay for a given job, so the earnings gap is due to differences in job composition across castes.

There are about 1050 elite jobs in my sample that participate in the on-campus job fair organized by the elite college. More than 95% of about 4200 college graduates participate in this job fair and submit applications for top MNCs. The on-campus job fair is highly structured in most elite Indian colleges with placement offices enforcing many rules that both firms and students must follow in order to participate (Section 4.4). *All jobs in my sample conduct four screening rounds, each of which is eliminatory.* In the first round, employers read applications and determine whether to advance students further. The second round is a written aptitude test. These tests, also called "technicals," are mostly conducted online. Technicals differ by job roles and usually include coding tests for manufacturing or technology jobs and case-study-based ("general aptitude") tests for consulting jobs.

The third screening round is a non-technical round called "group debate." Group debates are conducted for all jobs in my sample, whether client-facing or non-client-facing. These debates are typically conducted in multiple batches per job, with each batch evaluated by different members of the full recruiting team. A group debate batch typically comprises 30-35 students who are further divided into two teams of 15-20 students each. Group debates last a few hours and assess a wide array of socio-emotional skills, including: communication, mannerisms, consensus building, collegiality, confidence, and teamwork. Group debate topics are general in scope (e.g., "Is social media a boon or a bane?"). The fourth, and final, screening round is a non-technical personal interview.

Personal interviews are a widely employed discretionary screening practice by elite jobs hiring in India and elsewhere. Importantly, they are *not* technical or case-study-based and, are therefore, typically unstructured (see Section 6; Deshpande, 2011; Fernandez, 2018). In India, extensive survey evidence shows that shared experiences, educational qualifications of family members, neighborhood of residence, family background, father's job, cosmopolitan attitudes, upbringing, personal hobbies, and other aspects of "cultural fit" are common topics of conversation in such interviews (Deshpande and Newman, 2007; Jodkha, 2017).[2]

Free-form or unstructured interviews that screen on conversations related to aspects of cultural fit are not unique to India. Such practices are also common in the screening practices of elite U.S. colleges and corporate America (Stevens, 2009; Dana et al., 2013; Rivera, 2012, 2015). For example, Rivera (2012) shows that managers in the U.S. offices of 120 elite firms prioritized cultural fit in interviews largely because they believed, unlike job skills, "fit" could not be easily taught. Nearly 80% of the surveyed professionals endorsed the use of an interview evaluation heuristic called the "airplane test": would I like to be stranded

[2]In fact, such interviews are so commonplace in India that *Naukri.com*, India's leading urban job search portal with a market share over 60%, suggests that the "best way to answer this common interview question [when asked by recruiters to introduce oneself] is to tell the hiring manager about your education and family background" (Naukri.com, 2019).

in an airport with the candidate?

Recall that there was a 10% earnings gap across castes upon graduation from the elite college after controlling for pre-college skills, college academic performance, previous labor market experience, and other employer-relevant skills. In my setting, firms post uniform, job-specific (not match-specific) wages that are non-negotiable over the course of job search. Therefore, the relative contribution of each of the four screening rounds in explaining the caste earnings penalty of 10% among graduates can be quantified by comparing the caste difference in the composition of jobs that remain in contention at each job search stage. In other words, a measure of the "earnings" penalty can be constructed even before students formally realize their earnings, which is typically only after they accept job offers.

Almost all of the 10% caste earnings penalty occurs between personal interviews and job offers. Disadvantaged castes apply to similarly paying jobs as advantaged castes because a streamlined online application portal significantly reduces the marginal cost of an application. The three pre-personal interview screening rounds—application reading, written aptitude tests, and group debates that assess socio-emotional skills—together contribute to only about 10% of the earnings penalty among graduates. In addition, the composition of job choices over offered jobs contributes to a negligible portion of this earnings penalty. Therefore, when including controls, about 90% of the earnings penalty among graduates occurs between personal interviews and job offers (Figure 1).

The role of personal interviews in explaining caste disparities in hiring is much stronger if, instead, one were to quantify the relative contribution of each of the four screening rounds in explaining the lower success rate of disadvantaged castes compared to advantaged castes at the job offer stage. "Success rate" is defined as the probability that a given student makes it past a given screening round for a given job in a given year. Disadvantaged castes are about 30 percentage points (or about 115%) less likely to get job offers than advantaged castes. More than 95% of this caste penalty in the success rate can be attributed to the personal interview round (Table 6). Therefore, disadvantaged castes are almost as likely to succeed as advantaged castes at any of the three screening rounds prior to personal interviews. These results suggest that policies informing applicants about job opportunities, modifying student preferences, or improving performance at university are unlikely to meaningfully mitigate caste disparities.

Caste disparities in initial hiring outcomes are not aligned with productivity on the job, as disadvantaged castes often have better job performance once hired. While I do not observe direct measures of performance at any screening round, I tackle this limitation by collecting other, post-hiring measures of job performance. I collect these measures for the same jobs that were recruiting students at the elite college. Linking the sample to data from a major employment-focused platform provides two measures of job performance: tenure and promotion rates. Disadvantaged castes have about 15% higher job tenure than advantaged castes. Disadvantaged castes are also about 15% more likely to be promoted than advantaged castes. Similar results also hold among client-facing and non-client-facing jobs. Since I do not estimate promotion rates for the marginal candidate, I deal with potential issues due to inframarginality by constructing a measure of job selectivity. "Job selectivity" is proxied by calculating the cumulative proportion of job applicants from the elite college cut by a given job by the end of the final screening round comprising informal personal interviews. For example, if a job makes job offers to only 5% of its initial applicant pool, the cumulative proportion of applicants cut by the job by the end of the personal interview round is 95%. Using this measure, I categorize jobs according to their selectivity and quantify caste differentials in promotion rates among jobs with different levels of selectivity. Interestingly, disadvantaged castes also have about a 15%

higher promotion rate, *even among the most selective jobs* (Section 7.6).

The emergence of the caste penalty in initial hiring outcomes closely parallels caste revelation. Caste in elite, urban-educated India is almost perfectly signaled through caste-correlated characteristics such as family background, neighborhood of residence, and other aspects of "cultural fit" that are plausibly only revealed during personal interviews. Such characteristics are generally hard to fake because caste networks are deeply entrenched (Beteille, 1965, 1969; Jodkha, 2017). However, surface-level cues observed by employers during the earlier application reading, written aptitude test, and group debate rounds—such as last names, skin color, facial features, accents, and dialects—are highly noisy signals of caste in elite, urban-educated India, where there is large regional variation in such cues (Section 8). The salience of such cues is especially weak in settings such as elite colleges that are typically characterized by a lot of "mixing," as they draw students from myriad geographical regions.[3,4,5] Thus, the personal interview is likely the main mechanism through which caste is revealed, leading to direct or caste-based discrimination.

I also provide quantitative evidence that the caste penalty in hiring outcomes is plausibly linked to caste revelation. In a small subsample where last names provide clearer caste signals, the caste penalty emerges much earlier—during the application reading round. However, the *overall* caste penalty during the application reading round is negligible because most of the sample (about 85%) has non-distinctive caste names. Instead, the caste penalty primarily emerges during personal interviews, where students are screened based on background, hobbies, and cultural fit: characteristics strongly correlated with caste (Section 8).

Given the above evidence, my preferred interpretation of the caste penalty is that it stems from employers directly penalizing caste. Employers may penalize caste due to taste-based, statistical, customer, or coworker discrimination, unintentional biases due to homophily, or other factors. I do not distinguish between these sources of discrimination. Alternatively, one could also rationalize caste disparities as stemming from a scenario in which employers do not directly penalize caste. Under this scenario, employers could value other background characteristics that are correlated with caste. Consistent with survey responses of HR managers hiring for top MNCs in the U.S. and India, such background characteristics (e.g., family background, neighborhood of residence, and cultural fit) can plausibly only be learned during personal interviews and therefore cannot be used to differentiate candidates earlier (Section 6). Still, even when employers "only" value caste-correlated traits, but not caste directly, such screening practices could create barriers to economic mobility for disadvantaged groups, especially as impressions of cultural fit are often skewed by socioeconomic considerations. If the current equilibrium in India is such that students from disadvantaged castes do not have the opportunity to develop the "right" soft skills or are perceived to not hail from the "right" socioeconomic background, they might face significant barriers in accessing top jobs.

Although my preferred interpretation of caste disparities in hiring is that they stem from employers directly penalizing caste, the magnitude of the caste penalty is *invariant* to the true interpretation or source

[3] While many surnames are indicative of caste, most of them do not convey caste signals, especially in urban-educated settings due to migration and geographical variation in naming conventions. Indeed, in a recent audit study based on firms in the New Delhi area, Bertrand et al. (2009) state that the "enormous regional variations [in last names] mean that the precise coding of a particular last name is unlikely to be familiar to people from a different linguistic region of India." Surnames like "Singh," "Sinha," "Verma," "Chaudhary," "Mishra," and "Das" are shared across castes (Anthropological Survey of India, 2009). Moreover, naming conventions also differ significantly across regions. For example, for South Indians, personal (first) names often perform the role of traditional "surnames" (Jayaraman, 2005).

[4] Scholars have also argued that there is no association between skin color and caste, especially since Indian skin color is influenced mostly by geographic location rather than caste status (Mishra, 2015; Parameswaran and Cardoza, 2015).

[5] Perception of accent and dialect variation among young, English-speaking university graduates in India is linked to broad regional factors, such as a North Indian or South Indian accent, instead of caste status (Wiltshire, 2020).

of the caste penalty. Under a scenario where employers do not value caste at all, the "caste" coefficient in my regressions could be interpreted as a reduced-form capturing discrimination due to employers valuing other caste-correlated background characteristics learned by them during personal interviews. Caste disparities could stem from either taste-based or statistical discrimination and the reduced-form caste coefficient *would embed a mechanism for both.*

One potential alternative theory is that caste disparities might be concentrated at personal interviews because jobs potentially lower the standard and bump up weaker performers at earlier screening rounds to "appear fair" (ensure caste balance at each screening round), only to later use personal interviews to "catch up" to previous performance and not extend job offers to such students. Recall, I do not observe direct performance measures at any screening round. Instead, I only observe if students qualified past (i.e., survived) a given screening round. However, I can rule out a scenario where employers are lowering standards before personal interviews to appear fair. To do so, I first extend the previous definition of "job selectivity" to include the cumulative proportion of students cut by a given job by the end of *any* screening round. For example, if a job cuts 50% of its applicant pool at the application reading round and another 50% at the following written test round, then the cumulative proportion of applicants cut by the end of the written test round is 75%. I use large variation across jobs in the cumulative proportion of students cut by the end of each screening round to test if there is a substantial "caste advantage" (i.e., disadvantaged castes having a higher success rate compared to advantaged castes) before personal interviews, especially among the most selective jobs.

The intuition for this test is straightforward. Consider a job that cuts a high cumulative proportion of applicants before personal interviews. Such a job would probably have to lower standards considerably to achieve caste balance at each screening round, especially since disadvantaged castes have, on average, weaker academic credentials than advantaged castes. If there are enough such jobs in the sample, then conditional on characteristics such as GPA, one should expect disadvantaged castes to eventually have a *higher* success rate than advantaged castes before personal interviews, especially among the most selective jobs. Essentially, variation in job selectivity is used to truncate the distribution of latent candidate talent and examine if a caste advantage arises in the success rate when one goes further and further to the right of this distribution as candidate screening advances. I show that there is a negligible caste advantage in the success rate at any of the three screening rounds prior to personal interviews, essentially across the *entire distribution* of job selectivity. This result suggests that jobs are not lowering the standard to bump up weaker performers before personal interviews to appear fair. In other words, there appears to be no differential scrutiny by caste for the majority of the student sample before personal interviews. A similar argument also rules out caste differences in the propensity to be at the margins of the unobserved "cutoff thresholds" employed by jobs to advance candidates at screening rounds prior to personal interviews. Recall, a caste penalty in the success rate emerged immediately at the application reading round, but only amongst those with obvious caste names. If employers were trying to appear fair or virtue signal, one would think that students they would most want to give a bump to are those with obvious caste names.

Interview outcomes might also reflect inframarginal differences in performance at the screening rounds prior to personal interviews that could be used by employers to remove the worst *cumulative* performers among those who qualify for personal interviews. To rule out such a scenario, I again use large variation across jobs in the cumulative proportion of students cut by the end of each screening round. If there are large caste differences in performance before personal interviews, then disadvantaged castes should eventually have a *lower* success rate than advantaged castes at the screening rounds prior to personal interviews, especially among the most selective jobs. However, there is a negligible "caste disadvantage" (i.e., disadvantaged castes having a lower success rate than advantaged castes) before personal interviews among jobs with varying levels of selectivity. Only after personal interviews do disadvantaged castes have a notably lower success rate than advantaged castes, essentially across the *entire distribution* of job selectivity (Figures 4 and 5). These results are inconsistent with employers relying on cumulative performance differences at the screening rounds before personal interviews to advance disadvantaged castes at lower rates than advantaged castes after personal interviews. If cumulative performance differences before interviews were significant, one should have seen a caste disadvantage in the success rate at the three screening rounds prior to personal interviews, particularly among the most selective jobs.

Other alternative explanations are also unlikely to account for the caste penalty in initial hiring outcomes. These include differences in socio-emotional skills, outside options, negotiation abilities, competition from the government sector that has caste-based quotas, jobs making offers before conducting all four screening rounds, lack of meaningful student attrition at each screening round, and internal institutional pressure at elite MNCs to promote caste diversity (Section 10).

2 Contribution to the Literature

This paper makes several contributions, especially in labor and personnel economics. First, it presents unique descriptive evidence on the decomposition of hiring disparities that go beyond what one can learn from traditional resume-based correspondence studies. Unlike those studies, this paper accounts for rounds beyond the initial resume screening (first round), which are especially relevant since discrimination occurs during *later* stages in this setting. By studying the entire hiring process this paper shows that when it comes to caste and elite jobs, disparities in hiring almost exclusively arise in the final screening round—personal interviews—something that audit studies would not capture. Audit studies typically only measure group disparities in the first step of the hiring process, the evaluations of resumes or written applications. Thus, while audit studies measure something important, it is hard to say *how* important. Another well-known drawback of resume-based correspondence studies is that they typically cannot document effects on actual labor market outcomes, like job offers, job choices, or earnings (Bertrand and Duflo, 2016). To my knowledge, this is the first paper to link personal interviews, a common component of elite job recruitment in India and elsewhere, to actual labor market outcomes accumulated through the entire early career journey, including initial job offers, salaries, and even longer-term outcomes such as promotions.

Second, the paper's descriptive evidence advances the literature on the detection and measurement of labor market disparities. Disparities based on socioeconomic cues are likely to become more salient in elite, urban-educated settings because standard characteristics by which to differentiate groups become less perceptible as these settings become more multi-ethnic and diverse (Loury, 2002; Freeman et al., 2011; Gaddis, 2017). Therefore, collecting data on various stages of job search, each having the potential to reveal different socioeconomic cues, is crucial to better detect group disparities in the context of elite hiring. I provide an important example of the kinds of data future researchers may have to collect in such settings to better detect disparities from outwardly neutral screening practices. Even in much of the Western world, as racial boundaries blur, "caste" discrimination is likely to become a bigger part of the discourse globally (Austen-Smith and Fryer, 2005; Eguia, 2017; Kim and Loury, 2019; Rose, 2022). Indeed, as stated by

Rose (2022), group differences could often "emerge *endogenously* as an outcome of stereotypes that, in equilibrium, engender real disparities across groups" and therefore reinforce perceptions of racial, class, or caste differences. Discriminatory practices in elite jobs can often be coded through hiring managers valuing subjective aspects of fit, a practice that often favors socioeconomically advantaged groups. Yet, recent work studying the role of hiring discretion has either exclusively focused only on low-skilled jobs or provided correlational evidence between callback disparities and HR practices being more subjective (Hoffman et al., 2018; Kline et al., 2022). This paper blends insights from other fields such as sociology, psychology, and organizational decision making and presents a novel contribution by quantifying the role of personal interviews in explaining disparities in a real world, high-stakes setting. Many studies in the aforementioned fields have shown that unstructured interviews, evaluations, or other forms of subjective assessments are common avenues where disparities emerge (see e.g., Highhouse, 2008; Rivera, 2012; Dana et al., 2013). Therefore, the goal of this paper is not to posit unstructured, informal interviews as a "new" mechanism for disparities. However, this paper goes beyond showing that personal interviews are common avenues where disparities emerge, quantifying the real effects of disparities due to such subjective screening practices in earnings and access to elite jobs. To my knowledge, this paper is the first to quantify *how much* of the hiring disparities at elite jobs could be attributed to subjective assessments of fit done through personal interviews.

Third, the paper's descriptive evidence is helpful in identifying which policies are likely to be effective in remedying hiring disparities, in India and elsewhere. Potential policies need to be consistent with India's legal environment, which enforces explicit caste-based affirmative action policies, but only in public jobs and colleges. Therefore, unlike in the U.S., there is no federal ombudsman to regulate private-sector hiring. Moreover, while explicit caste-based discrimination is illegal, the Indian legal system does not recognize "disparate impact," and there is no systematic legal provision (anywhere) to penalize employers for judging "cultural fit" based on myriad characteristics correlated with protected status (Rivera, 2015; Jodkha, 2017; Lang and Spitzer, 2020). By quantifying the sources of caste disparities in elite jobs, policymakers could better understand the appropriate response in this context. Given that caste disparities primarily emerge after personal interviews in the multi-stage job search process, policies informing applicants about job opportunities, modifying student preferences, or improving performance at university are unlikely to meaningfully mitigate caste disparities. Therefore, the empirical quantification of the sources of caste disparities in elite hiring outcomes offers insights into better understanding the most effective policies to promote caste diversity among elite jobs hiring in India, nearly all of which hire entry-level employees from elite colleges such as the one in my setting (Section 4.2). Moreover, given potential productivity gains from hiring disadvantaged castes, elite multinationals hiring in India might even voluntarily adopt restrictions on the types of questions asked during informal interviews. Such policies also have relevance in the United States where, despite stronger employment laws, qualitative surveys have shown that the types of questions asked during open-ended, informal interviews tend to heavily favor the socioeconomically advantaged (Rivera, 2012, 2015). And even though the basis for hiring disparities in the United States is often "visual," reducing reliance on informal interviews could reduce socioeconomic and therefore racial or ethnic disparities in access to elite jobs. It is also not clear if hiring managers conducting informal evaluations have superior information or know what they are doing (Hoffman et al., 2018). While truly "blind" hiring à la Goldin and Rouse (2000) may not be possible in the context of elite job recruitment, reducing reliance on informal screening tools prone to bias is worth considering, in India and elsewhere.

Fourth, this paper contributes to a growing but small literature on the potential link between non-academic or "soft" characteristics used during candidate screening and eventual labor market outcomes. Chetty et al. (2023) show that non-academic characteristics (e.g., legacy preferences) used by elite U.S. colleges to screen candidates are weakly correlated with several measures of labor market performance. Benson et al. (2023) show that soft ratings on employee "potential" are often weakly correlated with job performance. Likewise, in the Indian context, Bhavnani and Lee (2021) show that personal interview scores used to screen candidates for top bureaucratic jobs are weakly correlated with performance. This paper complements these studies and examines the link between soft assessments of "fit" done through personal interviews as well as initial hiring outcomes and job performance among elite MNCs. Specifically, this paper quantifies how much of the initial hiring disparities among elite jobs could be attributed to personal interviews, a common component of elite job recruitment both in India and elsewhere. Further, by showing that disadvantaged castes have higher tenure and promotion rates, this paper suggests that disadvantaged castes are unwarrantedly held to higher scrutiny during such subjective evaluations.

Fifth, my empirical approach to capture the caste penalty through a reduced-form caste coefficient that potentially captures various sources of caste disparities helps advance recent research that argues for a constructivist understanding of group identities, instead of treating them as immutable facts (Hull et al., 2022; Rose, 2022). Such an approach is crucial to better understand "caste," classifications of which are rooted in the economic, political, and material history of India (Beteille, 1965, 1969). In addition, perceptions of caste in elite, urban-educated India are guided by a myriad of socioeconomic cues, paralleling the impressions of social class in other contexts, especially Britain and the United States (Deshpande, 2011; Mamidi, 2011; Savage, 2015; Jodkha, 2017).

3 Caste and Affirmative Action in India

Consistent with the practice of affirmative action policies in India, this paper focuses on two caste groups: advantaged ("upper") and disadvantaged ("lower") castes. This section provides a brief history of caste-based affirmative action policies in India and their present limitations. I will also emphasize that categorizations of "caste" have constantly interacted with social class and have been forged over political and historical processes spanning decades.

The first provisions for uplifting "depressed" or socioeconomically disadvantaged classes of Indian society were made possible after the Government of India Act of 1919 (also called the Montagu-Chelmsford Reforms) established self-governing institutions (i.e., provisional assemblies and central legislative assemblies), which introduced limited self-government to a majority British-controlled India. The electoral provisions under these reforms set aside seats for depressed classes in legislatures. The Government of India Act of 1935 replaced the words "depressed classes" with "Scheduled Castes" (Bayly, 2008). Many articles of the Constitution of India, ratified in 1949, formalized quota-based affirmative action policies in legislatures, higher-educational institutions, and government jobs for the so-called "backward" classes. Backward classes were intended to include not only members of Scheduled Castes (SCs) and Scheduled Tribes (STs), but also those from the Other Backward Classes (OBCs).[6] These provisions begged an obvious question: what determines "backwardness"?

In 1979, the Mandal Commission was set up with a mandate to "identify the socially and educationally

[6]Scheduled Castes (SCs), also known as Dalits, are socially marginalized in India. Scheduled Tribes (STs), also known as Adivasis or tribal groups, also face discrimination because of their ethnicity (Bayly, 2008).

backward classes in India," recommending caste as the basis for reservation (Mandal Commission Report, 1980). In particular, it recommended a 27% reservation (quota) in central and state services, public undertakings, and educational institutions for OBCs. Given the already existing 22.5% reservation for SCs and STs, the fraction of reserved seats for SCs, STs, and OBCs was brought up to 49.5%. The recommendations of the Mandal Commission were formally implemented in 1990. However, *none* of the current constitutional provisions extend to advancing affirmative action policies in the private-sector, where nearly all jobs in the sample of my study are located. Private-sector jobs also pay more than those in the public sector and thus are plausibly more relevant to the economic mobility of marginalized groups (Subramanian, 2019).

4 Key Definitions and Institutional Setting

I elaborate on the definitions of key terms used throughout the paper. I also describe the on-campus job fair of the elite college, the main institutional setting of this paper.

4.1 Disadvantaged Castes, Elite Jobs, and Elite Colleges

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

a. **Disadvantaged castes.** Consistent with government classification of caste categories, I define "disadvantaged castes" as those belonging to SCs, STs or OBCs (further defined in Section 3).

b. **Elite jobs.** An "elite entry-level job" is defined as a formal entry-level job in a multinational corporation (MNC), a major national corporation, or a unicorn firm hiring in the Indian private sector.[7] All jobs in my sample pay in the top 1% of entry-level job salaries in India (The State of Inequality in India Report, 2022). For the remainder of this paper, I use the terms "elite entry-level jobs" and "elite jobs" interchangeably. "Elite jobs" comprise 98% of all jobs in my sample.

   The remaining 2% entry-level jobs in my setting are in the public sector. Given the negligible presence of public-sector jobs in the sample, I condition the analysis on elite private-sector jobs and discuss the omission of public-sector jobs in Section 5.2.

c. **Elite colleges.** "Elite colleges" are those consistently ranked in the top 20 in their respective fields by *India Today*, which is the Indian equivalent of *U.S. News & World Report*. Examples of fields include science, engineering, arts and humanities, social sciences, and business. Examples of elite Indian colleges include the Indian Institutes of Technologies (IITs), Indian Institutes of Managements (IIMs), National Institutes of Technologies (NITs), and highly selective colleges under the umbrella of the prestigious Delhi University, such as St. Stephens, Sri Ram College of Commerce, and Lady Sri Ram College. I exclude elite Indian law and medical schools from my definition of "elite" colleges because the job search behavior of students attending such colleges is typically less industry-focused. Additionally, nearly all elite Indian colleges are public institutions (Altbach, 2012; Datta, 2017).

[7]Recall, a unicorn firm is typically a privately-owned startup with a valuation of over $1 billion.

4.2 Representativeness of the On-Campus Job Fair of the Elite College

The on-campus job fair of this elite college offers a representative window into how elite college graduates transition into elite jobs in India for four main reasons. First, almost all elite Indian colleges organize on-campus job fairs similar to the one studied in this paper. Second, almost the same group of firms recruits from on-campus job fairs across all elite Indian colleges. Third, elite entry-level jobs almost exclusively hire from elite colleges. Similarly, elite college graduates almost exclusively work in elite entry-level jobs. Fourth, all but a negligible proportion of college graduates search for jobs *solely* through the on-campus job placement fairs of elite Indian colleges (Online Appendix Section A.1).

4.3 The On-Campus Job Fair of the Elite College

The job fair of the elite college—the key institutional feature of the paper—takes place *entirely on the college's campus*. It can be divided into three broad phases: 1) pre-placement phase, 2) placement phase, and 3) post-placement phase or the "aftermarket."

a. **Pre-Placement phase.** The placement office invites firms prior to June. Between June to mid-August, firms visit the college campus and conduct pre-placement talks to advertise job profiles and gauge student interest. Firms make return offers from summer internships by late August. These offers are also called pre-placement offers (PPOs) and have late August deadlines. Students who accept their PPOs are *disallowed* ("de-registered") by the placement office from participating in the formal placement process for full-time jobs. Unlike in the U.S., return offers from junior year summer internships are *not an important feature* of full-time job search in elite Indian colleges. In my setting, less than 5% of the graduating cohort accepts return offers from summer internships and drops out of the regular, on-campus placement process for full-time jobs. Therefore, more than 95% of the graduating cohort participates in the on-campus job fair for full-time jobs and submits applications for top MNCs (Section 7.1.2).

   Participating students "register" for the formal on-campus job placement process by late August. By early September, firms submit *employer registration forms* to the placement office; these forms list details including job positions, compensation packages, and the probable number of slots or vacancies firms want to fill from the college that year. After these forms are submitted, advertised job profiles are considered "locked." They cannot be changed by firms during the course of the placement cycle. Moreover, students are prohibited by the placement office from bargaining over compensation bundles during the course of the placement process.[8] The placement office verifies advertised compensation bundles by requiring students to submit copies of their job offer letters and holds the power to prevent firms from returning to recruit on campus if the compensation package eventually offered to students deviates from what was advertised. It is beyond the scope of the paper to understand why certain placement rules exist. Such rules can be easily found on the websites of most elite Indian colleges.[9]

[8] Firms also appear to set salaries *nationally*, perhaps due to administrative convenience. The same firm will be touring multiple colleges for the same job profile and standardization of the compensation package for that job profile across all its Indian locations ("national wage setting") is likely a convenience measure. National wage setting has also been found in the labor market for elite entry-level jobs in the United States (Hazell et al., 2021).

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

b. **Placement phase.** Students start applying for jobs in mid-September. Jobs make the "first cut" after skimming through applications and invite students for additional screening. Next, jobs conduct written and verbal (socio-emotional skills) tests to determine eligibility for on-campus interviews. These three pre-personal interview screening rounds—application reading, written aptitude tests, and group debates—are conducted between September to early December. Jobs conduct the fourth and final screening round comprising subjective, non-technical personal interviews between December and January. After conducting personal interviews, jobs extend offers to students. Finally, students make job choices and the placement process concludes by January. *All jobs in my sample conduct all four screening rounds before making job offers*. About 70% of the entire student cohort secures full-time jobs by the end of the regular placement cycle in early January.

c. **Post-Placement phase or the "aftermarket."** Between late-January to the college's graduation ceremony in late-July, the placement office organizes an "aftermarket" for students who registered for the full-time placement process but did not secure jobs. These students comprise about 30% of the full student sample.[10] The recruiting process in this aftermarket relies heavily on referrals facilitated by the placement officer or other kinds of social networks. About 80% of the sample of students who participate in the aftermarket (about 25% of the full sample) eventually pursue other opportunities, such as entrepreneurship or advanced degrees (e.g., Master's and PhD), or take gap years to prepare for competitive exams for the civil services or elite MBA programs. The remaining 20% of students who participate in the aftermarket (about 5% of the full sample) find jobs in other firms.

   Firms that participate in the aftermarket are typically different from those that conduct recruitment only during the regular placement process for full-time jobs. The recruitment process in such firms is also much less standardized compared to the recruitment process at firms that hire only during the regular placement phase. As such, in this paper, I focus only on job recruitment that is conducted during the regular placement cycle for which the elite college's placement office collects data and students undergo a standardized hiring process comprising four screening rounds per job.

   In Section 7.1.2, I discuss the potential impact on my main results due to the selection induced by those students who do not participate in the regular placement process for full-time jobs either because they accept summer internship offers or are initially unplaced and therefore participate in the "aftermarket."

4.4 Key Rules of the College's On-Campus Job Fair

As mentioned in Section 4.2, almost all other elite Indian colleges organize on-campus job fairs similar to the one studied in this paper. Elite Indian colleges typically also set similar rules that both firms and students must follow to participate in the on-campus job fairs of such colleges. As mentioned in Section 4.3, it is beyond the scope of the paper to understand why certain placement rules exist. Such rules can be easily found on the websites of most elite Indian colleges.[11] The following are some key rules of the on-campus job fair set by this elite college's placement office for students and firms that participate in the regular placement phase for full-time jobs:

[10] Most such students tend to go on the market "softly" either because they already plan on pursuing other opportunities outside of finding employment in a select few firms or because they cast a narrow net while initially searching for jobs. Therefore, the trend among such students is that they either apply to very selective firms or startups in niche fields.

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

a. **Interview day allotment.** Each firm is allotted one "interview day" by the college's placement office to conduct the final screening round comprising personal interviews. There are typically between 7-10 interview days, spread over 2-3 weeks. Unlike job recruitment at U.S. colleges, there are *no further onsite interviews*. A particular rule of the on-campus job fair is that conditional on getting a job offer on a given interview day, a student can no longer participate in interviews on future interview days. At best, a student can receive multiple job offers within a given interview day. If a student does not get any job offer on a particular interview day, he can participate in interviews on future interview days.[12]

b. **Students cannot "reject" firms midway or accept job offers "early."** Having applied to a job, students are prohibited by the placement office to skip any of three rounds of screening conducted by that job prior to personal interviews. Students are also not allowed by the placement office to skip the sequence of scheduled job interviews that are typically spread over multiple interview days. Students are also not allowed to accept offers "early" in the process (e.g., by negotiating offers midway before they are officially announced for others).

c. **All job offers are announced at the end of the interview day.** All job offers are announced within a short interval of time at the end of the interview day to prevent jobs that are allotted the same interview day from coordinating on offers.

5 Data

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

5.1 Students, Admissions Criterion, and Baseline Characteristics

5.1.1 Broad Overview

This section provides a broad overview of the span of the data, the composition of the student body, and the admissions criterion for the elite college.

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

2. **The caste share within each college major is nearly equal.** Table 1 shows the total number of students belonging to each caste within each college degree. There are a total of 4164 students in the sample that are about evenly split between castes. Male students comprise about 90% of the sample (not shown).[13]

[12] Salaries are typically one of the most important determinants of interview day allotment. However, firms do not appear to strategically respond to such interview day allotment by altering their compensation package to be seen as more "desirable," so that they could get an earlier interview slot. For example, as mentioned in Footnote 8, job salaries appear to be set *nationally*, perhaps due to administrative convenience. The same firm will be touring multiple colleges for the same job profile and standardization of the compensation package for that job profile across all its Indian locations ("national wage setting") is likely a convenience measure. National wage setting has also been found in the labor market for elite entry-level jobs in the United States (Hazell et al., 2021).

[13] The fraction of males in the data is typical of those at elite technical colleges in India (Datta, 2017).

Caste in the data is reported by students, typically by presenting a formal "caste certificate" to the college and is used as the basis for quota-based policies in admissions. These policies equalize the share of both disadvantaged and advantaged castes *within each college major*. Therefore, nearly 50% of the students are from disadvantaged castes for three of the four college degrees in the sample (Section 3).[14]

3. **There are four college degrees in the sample.** Table 1 shows that there are four college degrees in the sample. They are the Bachelor of Technology (B.Tech.), "Dual," Master of Technology (M.Tech.), and Master of Science (M.S.) degrees. A Dual degree integrates undergraduate and post-graduate studies and is completed a year after the conventional four-year degree (B.Tech.). I omit students pursuing a different Master of Science (M.Sc.) degree, as they comprise less than 2% of the student population.[15]

4. **National-level entrance exam scores are the only criterion for college admission.** Students are admitted to the elite college through national-level entrance exams based on caste-major-specific cutoffs. These exam score cutoffs comprise the *only* criterion for college admission.

Table 1: Total Students by Caste and Degree

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

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

[14] In my sample, the Master of Science (M.S.) degree has a slightly larger proportion of advantaged castes. Frisancho Robles and Krishna (2015) also document similar patterns indicating that despite admissions quotas, some college degrees may not be able to fill all reserved seats.

[15] There are two different Master of Science degrees offered at the college. They are M.S. and M.Sc. The first, which is included in the sample, has an industry focus. The second, which is not included in the sample, is a much smaller program with a research or academic focus. Students pursuing an M.Sc. degree do not typically search for jobs through the college's placement office.

5.1.2 Student Characteristics

This section discusses caste differences in some key student characteristics.

1. **College GPA.** The data includes the college GPA of students. There are substantial caste differences in college GPA within each degree (Table 2). Importantly, there are students from both castes (i.e., common support) within each GPA decile (Online Appendix Figure OA.1).

2. **Pre-college skills.** The data includes entrance exam (EE) scores of students. As mentioned in Section 5.1.1, EE scores are the only criterion for college admission. EE scores in the data are *originally exam ranks* that have been renormalized so that larger numbers are better. Caste differences in EE scores are large (Table 2). As with college GPA, there are students from both castes within each EE score decile (not shown).

3. **Previous labor market experience.** The data also include detailed information on both summer and winter internships, including duration of internship employment, duration of part-time or full-time employment, total pay during internships, total pay during part-time or full-time employment, sector of internship employment, and employment in startups. Internship descriptions typically include application eligibility criterion and desired skills. Application eligibility depends upon a combination of factors such as major, degree, and GPA. For example, a software engineering position at Microsoft will typically restrict applications to students from the Computer Science major. An analyst position at Goldman Sachs will typically invite applications from all college majors.

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

   The weak correlation between internship outcomes and pre-college skills is a common pattern in elite Indian colleges. Unlike in the U.S., return offers from junior year summer internships are *not an important feature* of full-time job search in elite Indian colleges. Indeed, as mentioned in Section 4.3, less than 5% of students accept return offers from summer internships and skip the regular placement fair of this elite college. Similar proportions of students skip the job fairs of other elite Indian colleges (Online Appendix Section A.1).

   Informal conversations with employers and students also suggest that internships are less of a matching problem per se. Students seem to care more about the job sector or the firm instead of salaries; apart from a select few firms, students are just happy to get any internship and vice-versa. Therefore, internships are less of a matching problem, which could also explain modest caste differences in internship outcomes.

4. **Other employer-relevant skills.** These measures include college major, college degree, and coursework. Admissions quotas coupled with fairly rigid college curricula lead to almost no caste differences on many measures of these employer-relevant skills. For example, as mentioned in Section 5.1.1, the

caste share within each college major is roughly 50% because of quota-based admissions policies within each major. Similarly, the caste share within each college degree is about 50%.

Table 2: Summary Statistics (Students)

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

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

5.2 Definition of a Job and an Overview of Jobs Recruiting from the Elite College

This section provides the definition of a "job" and an overview of jobs recruiting from the elite Indian college.

1. **Definition of a job.** A "job" is defined as a job designation within a firm. For example, Google can hire a product manager and a software engineer. These are two different jobs.

   Over 98% of the jobs in my sample are in the Indian offices of firms. When hiring for Indian job positions, employers conduct screening rounds for all Indian offices at once. The specific Indian office that a candidate ultimately gets assigned to is decided later and usually depends on the location-specific needs of that employer. For example, Google will conduct a combined recruitment process for its software engineering role regardless of whether the candidate may ultimately get assigned to either the Mumbai or Bangalore office of Google. Therefore, in this paper, a "job" will almost always mean a firm and job designation pair and not the ultimate (Indian) office location. Moreover, as mentioned in Section 4.3, the compensation package for a given employer and job designation pair is typically standardized across all its Indian locations.[16] The standardization of compensation packages is also compatible with how recruitment is conducted for the Indian offices of employers.

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

2. **Location of jobs.** The majority of elite jobs are within multinational corporations headquartered in the U.S. or Europe. Many of these firms are regularly featured in the Fortune 500. Although the majority of the jobs are based in firms headquartered outside India, almost all jobs in my sample hire for their Indian locations (Table 3).

3. **Number of jobs in the sample.** There are about 1050 jobs in the sample. Only about 2% of the jobs in the sample are in the Indian public sector (Table 3). The rest are in the private sector.

4. **Number of years each job recruits from campus.** As mentioned in Section 5.1, the data belong to four placement cycles or years. While there are four full cohorts of students in the sample, each job may not recruit from the college campus during each of the four placement years. Some jobs recruit

[16] As mentioned in Footnote 8, job salaries appear to be set *nationally*, perhaps due to administrative convenience. The same firm will be touring multiple colleges for the same job profile and standardization of the compensation package for that job profile across all its Indian locations ("national wage setting") is likely a convenience measure. National wage setting has also been found in the labor market for elite entry-level jobs in the United States (Hazell et al., 2021).

during only one placement cycle while others recruit from campus during each of the four placement cycles. Most jobs in the sample recruit from this college's campus for at least two out of the four placement cycles.

5. **Almost all hiring professionals are Indians.** As mentioned above, almost all jobs in my sample hire for their Indian offices. Therefore, while I do not have data to quantify recruiter nationality, it is reasonable to assume that recruitment is typically conducted by Indian hiring professionals or managers.

6. **Omitting public-sector jobs.** I omit public-sector jobs from my analysis for two reasons. First, such jobs comprise less than 2% of all jobs in the sample (Table 3). Second, public-sector jobs in India are quite different from their private-sector counterparts, especially in areas like salary structure and job stability.[17]

7. **Distribution of jobs and salaries by job type.** Table 3 shows the distribution of jobs and the average salary across all jobs by job type. About 76% of jobs are non-client-facing and the remaining 24% are client-facing.

   Jobs set uniform, job-specific (not match-specific) salaries. Conditional on college degree, job salaries do *not* vary across major, caste, or gender. Average salaries across all jobs in client-facing and non-client-facing jobs are about $57,553 (PPP) and $61,116 (PPP), respectively. The average salary across all jobs is about $60,265 (PPP).

8. **Client-facing and non-client-facing jobs.** As shown in Table 3, non-client-facing jobs comprise almost 75% of all available jobs.[18] Almost all jobs in the technology sector are non-client-facing whereas almost all jobs in the consulting sector are client-facing. The manufacturing sector has a roughly even mix of client-facing ("Management Trainee" and "Manager") and non-client-facing ("Graduate Engineer" and "Structural Engineer") jobs.

9. **Job descriptions and non-pecuniary amenities.** Jobs declare job details in the *employer registration forms* made available to them by the college's placement office. These details include information on various job characteristics, including job descriptions, job designations, sector, salaries, non-pecuniary amenities, desired skills, expectations on the job, the expected number of slots or vacancies a job wants to fill from the college, and even job application eligibility criterion (see "Pre-Placement Phase" in Section 4.3). Application eligibility typically depends upon a combination of factors such as major, degree, and GPA (see point 3 in Section 5.1.2). The number of advertised non-pecuniary amenities is high and ranges between 40-50 per job.

10. **Return offers from internships.** As mentioned in Section 4.3, firms make return offers from summer internships by late August. These offers are also called PPOs or pre-placement offers and have late August deadlines. Students who accept their PPOs are disallowed ("de-registered") by the placement office from participating in the formal placement process for full-time jobs. Therefore, a typical

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

[18] Detailed job descriptions (particularly, job titles and job functions) were used to categorize jobs as "client-facing" versus "non-client-facing." Typically, a software engineering role would be considered non-client-facing, whereas a consulting or managerial role would be considered client-facing.

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

Unlike in the U.S., return offers from junior year summer internships are *not an important feature* of full-time job search at elite Indian colleges (see Section 4.3). In my setting, less than 5% of the graduating cohort accepts return offers from summer internships and drops out of the regular, on-campus placement process for full-time jobs. Therefore, more than 95% of the graduating cohort participates in the on-campus job fair for full-time jobs and submits applications for top MNCs (Section 7.1.2).[19]

Table 3: Summary Statistics (Jobs)

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

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

5.3 Data on Job Selectivity and Screening Rounds

This section shows that there is substantial variation across jobs in the proportion of the candidate pool cut both *at* as well as *by the end of* each screening round. It also provides more details on each screening round and the available data from it.

5.3.1 There is substantial variation across jobs in the proportion of the candidate pool cut both at as well as by the end of each screening round

All jobs conduct four rounds of screening before making job offers. Each screening round is eliminatory. There is also substantial variation across jobs in the cumulative proportion of students cut both at as well as by the end of each screening round. Table 4 shows the distribution of both the conditional and cumulative cuts across jobs. A "conditional cut" is the proportion of candidates cut by a given job at a particular screening round. A "cumulative cut" is the total proportion of candidates cut by a given job by the end of a particular screening round. For example, if a job cuts 50% of the candidates at each of the first two screening rounds, then the conditional cut made by the job at each round is 50%, whereas the cumulative cut

[19] I discuss the potential impact on my main results due to the selection induced by this placement rule regarding return offers from internships in Section 7.1.2.

made by the end of the second round is 75%. As shown in Table 4, students have to almost always be within the top 5% of a given job's initial applicant pool to get offers. It is also common for a decent proportion of jobs to not make any offers, as job search sometimes fails for a variety of reasons.

Table 4: Variation in Conditional and Cumulative Selectivity Across Jobs

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

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

5.3.2 Timeline of job search and details on the four screening rounds conducted by each job

The data comprise the list of job applications, written aptitude tests, group debates, interviews, job offers, and final job choices for each student. Students apply for jobs through a centralized job application portal, similar to Job Openings for Economists. Applying to a job only involves "clicking" on the name of the job in the online job application portal and does not require additional cover letters or other statements. Employers only request student resumes that are automatically made available to them when the student clicks on the name of a job to apply. The streamlined job application process significantly reduces the marginal cost of an application. Application eligibility depends upon a combination of factors such as major, degree, and GPA—information that the dataset contains (see also point 3 in Section 5.1.2).

Timeline of Job Search. As mentioned in Section 4.3, job search typically begins in September when students submit applications for full-time jobs. After applications are submitted, firms conduct written tests and group debates from late September to late November. The final screening round, comprising on-campus personal interviews, is typically conducted between December to January.

Four Screening Rounds Per Job. Each job in my sample conducts four, eliminatory screening rounds before making job offers. The four job screening rounds are as follows:

1. **Application reading.** In the first screening round, employers read applications and determine whether to advance students further. This "first cut" is typically made using a combination of college GPA, major, degree, and other aspects on student CVs (e.g., summer internships and exceptional performance

in international competitions). Informal conversations with employers suggest that college GPA is one of the main filters used during the application reading round. As mentioned in Section 4.3, the first screening round is usually conducted in September after applications are submitted.

Following the application reading round, the second and third screening rounds are conducted between September to late November (see also "Placement phase" in Section 4.3).

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

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

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

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

      During group debates, recruiters make students discuss a general or non-technical topic with one team speaking for each side. Meanwhile, recruiters behave like passive observers, with their only active role consisting of duties such as starting and ending the discussion on time and ensuring decorum. Discussion moderators organically emerge from the group of students participating in the debate.

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

4. **Non-Technical Personal interview.** The fourth and final screening round is a subjective, non-technical personal interview. Section 6 below provides qualitative evidence on the content of non-technical personal interviews conducted by recruiters at elite MNCs hiring both in India and the United States. As mentioned in Section 4.3, the final screening round comprising non-technical personal interviews is usually conducted between December and January.

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

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

c. **Almost all interviewers are likely Indians.** Recall, I do not directly have information on interviewer nationality. However, given that almost all firms in my sample, including foreign-based MNCs, hire for their Indian locations, it is reasonable to assume that almost all interviewers are likely Indians (Section 5.2).

d. **Interviews are almost always unstructured.** As will be elaborated in Section 6 below, interviews are usually unstructured. Indeed, a key emphasis in such interviews is to largely assess fit and other soft aspects. As such, these interviews are typically non-technical, discretionary, and therefore, unstructured.

5.4 Summary Statistics on Candidate Screening

This section discusses some summary statistics on candidate screening. The average (median) number of applications per job is 208.85 (140.17). Similarly, the average (median) number of written tests conducted, group debates conducted, interviews conducted, and job offers extended per job are 124.91 (87.00), 101.54 (70.00), 26.07 (18.00), and 6.49 (4.00), respectively.

Recall, application eligibility depends upon a combination of factors such as major, degree, and GPA. For example, a software engineering position at Microsoft will typically restrict applications to students from the Computer Science major. On the other hand, an analyst position at Goldman Sachs will typically invite applications from all college majors (see point 3. in Section 5.1.2). Such application eligibility criteria are typically why, despite the low marginal cost of an application, a student cannot apply to all jobs recruiting from campus in a given year (see Section 5.3.2).

The summary statistics on candidate screening suggest that jobs are typically very selective. On average, they conduct about 25 interviews and make only about 6 offers out of the average applicant pool comprising about 200 students. Therefore, to get an interview, one typically has to be within the top 10% of the initial applicant pool. Similarly, to get an offer, one typically has to be within the top 5% of the initial applicant pool. The summary statistics discussed in this section are also consistent with the variation across jobs in "cumulative cuts" (defined in Section 5.3.1) shown in Table 4.

Table 5: Summary Statistics (Candidate Screening)

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

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

5.5 Data on Job Tenure and Promotions

Linking the sample to data from a major employment-focused platform provides two measures of job performance: tenure and promotion rates. I can link over 98% of the sample of students who received jobs through the job fair with post-hiring measures of job performance. Job tenure is measured in months. "Promotion" is defined as a movement to a job designation higher in the firm hierarchy. To better capture "true" promotions, I compare job designations with job tenure. For example, an "Associate" position at a given consulting firm typically requires 2 years of experience. Thus, a student moving from an "Analyst" to an "Associate" position within that consulting firm would be considered as having been promoted. Promotions are a particularly good measure in this setting because entry-level jobs at elite multinationals typically have rigid salary structures, in India and elsewhere (Hazell et al., 2021; Sarsons et al., 2022).

5.6 Data Limitations

There are a few salient data limitations. First, I do not have access to student resumes. Student resumes are written in a standardized format prescribed by the placement office. I also do not have access to performance measures at any screening round. I address potential limitations due to lack of access to such data and their implications for my findings in Sections 9 and 10.

6 Qualitative Evidence on the Content of Personal Interviews at Elite MNCs Hiring in India and the United States

This section presents qualitative evidence on the content of personal interviews at elite MNCs hiring in India and the United States. Personal interviews are typically *not* technical or case-study-based. There are three key takeaways from this section. First, non-technical personal interviews are highly discretionary. Second, such interviews are avenues where informal conversations about shared experiences, personal hobbies, educational qualifications of family members, father's job, neighborhood of residence, preference for living in a cosmopolitan city, desire for traveling, and other aspects of "cultural fit" are considered par for the course. Third, such interviews are commonly used in the recruitment practices of elite MNCs hiring both in India and the United States.

This section should not be interpreted as providing dispositive evidence on the content of personal interviews in my setting. I do not have access to interview transcripts. Nevertheless, it is important to emphasize that extensive survey evidence in high-stakes settings, both in India and the U.S., has shown that unstructured interviews are highly discretionary and serve as common avenues for soft assessments of culture or "fit." One key contribution of this paper is to quantify the importance of personal interviews, a common component of elite job recruitment in India and elsewhere, in explaining unwarranted hiring disparities across socioeconomic groups (caste).

6.1 Personal interviews at elite MNCs hiring in India

In a detailed survey of HR managers at elite MNCs in the Indian private sector, Jodhka and Newman (2007) documented opinions regarding the role of personal interviews in the hiring process. These firms employ a total of nearly 2 million workers. The authors found unanimous agreement on the importance of asking questions pertaining to background characteristics in personal interviews. Interestingly, HR managers saw little contradiction in judging a candidate's individual merit through background characteristics, arguing instead that trustworthiness to potential clients, attrition, and cultural fit can be identified through such inquiries. One HR manager expressed what he looks for in personal interviews:

> We also ask a lot of questions related to family background. Questions like how many family members are there, how many are educated, etc. The basic assumption behind these questions is that a good person comes from a good and educated family. If parents have good education, the children also have good education. Some questions about their schooling . . . and the locality where they [grew up].

Another HR manager expressed the signaling value of background characteristics by mentioning that conversations about family background and upbringing are useful in forming impressions about trustworthiness to potential clients, cultural fit, attrition, and long-term professional behavior:

> As personal traits are developed with the kind of interaction you have with society . . . Where you have been brought up, the kind of environment you had in your family, home, colony and village, these things shape the personal attributes of people . . . This determines his behavior, and [his ability to work] in a group with different kinds of people. We have some projects abroad, and if a person doesn't behave properly with them, there is a loss for the company. Here the family comes in, whether the person behaves well and expresses himself in a professional way, for a longer term and not for a short term . . . This is beneficial.

Relatedly, in a separate survey of students from elite colleges in New Delhi, Deshpande and Newman (2007) found that almost all of them were asked about their family background and upbringing in personal interviews. Personal interview questions also included queries into their hobbies and cosmopolitan attitudes, like preference for living in a big city. As one student put it:

> Most of my interviews were very relaxed. No one was assessing my knowledge or anything, but . . . seeing how well and efficiently I contribute to the company . . . For example, when I had my interview with [information firm], he asked me why I want to work in Bombay? . . . So the interview was more in terms of what I like, what I dislike and general chit chat about what I was looking to do in the future rather than quizzing me about, let's say what particular topics I had done in a particular [academic] subject or something like that.

These prevailing personal interview conventions have also influenced the advice provided to candidates by major Indian job portals. *Naukri.com*, India's largest urban job search portal with over 60% market share, suggests that the "best way to answer this common interview question [when asked to introduce oneself] is to tell the manager about your education and family background" (Naukri.com, 2019).

Quantitative research on Indian labor markets has also supported the above qualitative evidence. Using Likert scale ratings, Mamgain (2019) shows that HR managers at elite firms value family background as much as candidate experience, quality of the institution of education, aptitude, or technical skills.

6.2 Personal interviews at elite MNCs hiring in the United States

Free-form or unstructured interviews that screen on conversations related to shared experiences, cultural fit, and personal hobbies are not unique to India. Such practices are also common in the screening practices of elite U.S. colleges and corporate America (Stevens, 2009; Dana et al., 2013; Rivera, 2012, 2015). For example, Rivera (2012) shows that managers in the U.S. offices of 120 elite MNCs prioritized cultural fit in interviews largely because they believed, unlike job skills, "fit" could *not* be easily taught. One evaluator categorically desired biographic similarities:

> I usually start an interview by saying, "Tell me about yourself." When I get asked that, I talk about where I'm from, where I was raised, and then my background . . . I want to hear your life story. Hopefully there's something more interesting about your life than deciding to go to school . . . When they tell me about their background, it's easier to find things in common . . . Maybe . . . they're from Seattle and I've been to Seattle. We can talk about that and develop a connection.

Furthermore, nearly 80% of the professionals surveyed in Rivera (2012) endorsed the use of an interview evaluation heuristic called the "airplane test." As one manager explains:

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

7 Descriptive Evidence

I document large earnings disparities across castes, argue that the earnings gap is conservative, and show that caste disparities in hiring outcomes arise primarily due to the final round comprising subjective, non-technical personal interviews. I also show that caste disparities are concentrated following personal interviews in a substantial proportion of jobs in my sample and discuss heterogeneity in my results by job type. Finally, I discuss caste differences in job tenure and promotion rates.

7.1 Large Earnings Gap Across Castes

About 70% of the entire student cohort secures full-time jobs by the end of the regular placement cycle in early January. Recall that less than 5% of the full student sample deregisters from the job placement process for full-time jobs by accepting summer internship offers. Therefore, more than 95% of the graduating cohort participates in the on-campus job fair of the elite college and submits applications for top MNCs. Out of this sample of participants, about 2550 students or 65% get placed into jobs by participating in the on-campus job fair. Students who register for the job fair but remain initially unplaced (about 30% of the full sample) typically participate in the "aftermarket," where recruitment relies on referrals facilitated by the placement officer.[20]

As mentioned in Section 4.3, firms that participate in the aftermarket are typically different than those that participate in the regular placement process for full-time jobs. In addition, relative to recruitment during the regular placement cycle, recruitment during the aftermarket is much less standardized both across firms and across jobs within firms (Section 4.3). About 80% of the sample of students who participate in the aftermarket (about 25% of the full sample) eventually pursue other opportunities, such as entrepreneurship or advanced degrees (e.g., Master's and PhD), or take gap years to prepare for competitive exams for the civil services or elite MBA programs. The remaining 20% of students who participate in the aftermarket (about 5% of the full sample) find jobs in other firms.

The primary focus of the paper is to study whether informal, non-technical personal interviews are avenues where employers (unwarrantedly) put a higher threshold on disadvantaged castes through subjective assessments of "fit." Thus, I focus the discussion of the descriptive evidence in Section 7.1 only on the sample of students who obtained full-time jobs through the college's on-campus job fair, and therefore underwent a standardized hiring process comprising four screening rounds per job.[21] As discussed in Section 7.1.1 below, nearly every student who secures employment through the job fair starts in the same position after graduation. Therefore, data provided by the placement office provide an accurate representation of a student's full-time employment status once he leaves college.

Throughout the paper, "caste earnings gap" is defined as the average earnings of disadvantaged castes minus the average earnings of advantaged castes. The earnings gap across castes among graduates from the elite college is -0.231 (0.020) log points or about 23% without controls, where the number in parentheses denotes the standard error. In the presence of detailed controls for pre-college skills, college academic performance, previous labor market experience, and other employer-relevant skills, the remaining gap is -0.103 (0.017) log points, or about 10% (Online Appendix Table OA.1). Relative to the mean earnings of graduates, this gap is about \$6000 (PPP). I discuss heterogeneity in these results by job type in Section 7.5.

7.1.1 Matching with Exit Data to Check the Accuracy of Job Status and Salaries After Graduation

"Exit data" is a college-administered survey of students conducted two months after graduation, that includes their job designations and responses to whether offer terms were negotiated between the conclusion of the regular placement process (around early January) and the rollout of the exit survey. Exit surveys are rolled

[20]Recall, most such students tend to go on the market "softly" either because they already plan on pursuing other opportunities outside of finding employment in a select few firms or because they cast a narrow net while initially searching for jobs. Therefore, the trend among such students is that they either apply to very selective firms or startups in niche fields. See also Footnote 10.

[21]In Section 7.1.2, I discuss the implications on the results reported in Section 7.1 due to the selection induced by those who do not participate in the regular placement cycle either because they accept summer internship offers after their junior year or are initially unplaced and therefore typically participate in the "aftermarket" (see Section 4.3).

out close to graduation (between May and June) and before the job start date (around late July). The data also include specifics of the negotiated terms and the reasons for negotiation. Note that students are restricted from bargaining over advertised compensation bundles only *during* the course of the regular placement process and not after its conclusion (see "Pre-Placement phase" in Section 4.3). The placement office administers exit surveys to validate its employment data.[22]

Exit data is filled by nearly 98% of the graduating cohort, possibly since passing along information regarding their current roles and contact details is relevant for students to take advantage of the elite college's extensive alumni network. Similar exit surveys collected by the career offices of elite U.S. MBA programs have comparable response rates (Exit Survey of MBA Graduates, Yale School of Management, 2021). Reassuringly, exit data show that basically all job getters began in the same jobs they obtained through the placement process and did not negotiate salaries or non-pecuniary amenities, even two months after graduation or several months after the conclusion of the regular placement process.

7.1.2 Selection and Implications for the Earnings Gap

I discuss differential selection by caste among those omitted from the earnings regressions discussed in Section 7.1.

1. **Less than 5% of the student sample "deregisters" from the job placement process of this elite college.** "Deregistered" students are those who either do not sign up to participate in the job placement process for full-time jobs or accept their return offers from summer internships in their junior year, and therefore are disallowed by the placement office to participate in the job placement process for full-time jobs (see "Pre-Placement phase" in Section 4.3).

   Less than 5% of students deregister from the on-campus placement process of the elite college. This number can be directly seen by comparing the total number of students in the sample with the total number of students who ultimately participate in the job fair (i.e., "register") and apply for full-time jobs. Out of a total of 4164 students in the sample, 4005 students apply for jobs recruiting from the on-campus job fair of the elite college (Table 1 and Online Appendix Table OA.2). The remaining 159 students are "deregistered" either because they accepted their summer internship offers or because they wanted to pursue other opportunities. Almost all deregistrations before the beginning of the regular placement process are due to the acceptance of junior year summer internships. Similar proportions of students skip the job fairs of other elite Indian colleges (Online Appendix Section A.1). In summary, more than 95% of about 4200 graduating students participate in the job fair of the elite college and submit applications for top MNCs.

2. **The earnings gap is conservative.** In addition to belonging to the "deregistered" category, students omitted from the earnings regressions reported in Section 7.1 could also belong to the "registered" category. In principle, omitted students from the registered category could comprise those who participate in the full-time job placement process but either do not get full-time jobs or reject all of their job offers. In practice, no one who registers for the formal placement process rejects all of their job offers, so this additional omitted category typically comprises only those who fail to receive any job offer during the regular placement phase of the on-campus job fair.

[22] To further validate exit data, I also match full-time job status of students with information obtained from a leading employment-focused platform (see also Section 7.6).

I compare the GPAs of students omitted from the earnings regressions with the GPAs of those who got jobs through the college's job fair. I find that, relative to the sample of students who got jobs, disadvantaged castes omitted from the earnings regressions are much *more negatively selected* on college GPA than advantaged castes (Online Appendix Table OA.3). This result is intuitive because students who typically do not get jobs (i.e., the majority of students omitted from the earnings regressions) are negatively selected on college GPA and most of the extremely negatively selected students on college GPA belong to disadvantaged castes. Since average earnings are increasing in GPA (not shown), the reported earnings gap is conservative.

7.2 Almost All of the Caste Earnings Penalty Is Due to Subjective, Non-Technical Personal Interviews That Often Assess "Fit"

This section lays out one of the key contributions of the paper. It presents unique descriptive evidence on the decomposition of hiring disparities that go beyond what one can learn from traditional resume-based correspondence studies. Unlike those studies, this paper accounts for rounds beyond the initial resume screening (first round), which are especially relevant since discrimination occurs during *later* stages in this setting. By studying the entire entry-level hiring process among elite jobs, I show that disparities in hiring almost exclusively arise in the final screening round—personal interviews—something that audit studies would not capture. Audit studies typically only measure group disparities in the first step of the hiring process, the evaluations of resumes or written applications. Thus, while audit studies measure something important, it is hard to say *how* important. Another well-known drawback of resume-based correspondence studies is that they typically cannot document effects on actual labor market outcomes, like job offers, job choices, or earnings (Bertrand and Duflo, 2016). To my knowledge, this is the first paper to link personal interviews, a common component of elite job recruitment in India and elsewhere, to actual labor market outcomes accumulated through the entire early career journey, including initial job offers, salaries, and even longer-term outcomes such as promotions (see e.g., Section 7.6).

Recall, firms post uniform, job-specific (not match-specific) wages that are non-negotiable over the course of job search (see "Pre-Placement phase" in Section 4.3). Therefore, the relative contribution of each of the four screening rounds in explaining the caste earnings penalty of 10% among graduates can be quantified by comparing the caste difference in the composition of jobs that remain in contention at each job search stage. In other words, a measure of the "earnings" penalty can be constructed even before students formally realize their earnings, which is typically only after they accept job offers.

I show that almost all of the caste earnings penalty of 10% among graduates occurs due to the final round comprising subjective, non-technical personal interviews. To do so, I run different specifications of the following linear regression:

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

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

For each student i, "average job salary" at a given job search stage s is calculated by averaging over the posted salaries of jobs j that remain in contention for that student at that job search stage. The coefficient of interest is β, which is shown in Figure 1 for each job search stage. Online Appendix Table OA.2 reports these regressions.

Figure 1 shows that almost all of the 10% caste earnings penalty occurs between personal interviews and job offers. Disadvantaged castes apply to similarly paying jobs as advantaged castes because a streamlined online application portal significantly reduces the marginal cost of an application. The three pre-personal interview screening rounds—application reading, written aptitude tests, and group debates that assess socio-emotional skills—together contribute to only about 10% of the earnings penalty among graduates. In addition, the composition of job choices over offered jobs contributes to a negligible portion of this earnings penalty. Therefore, with controls, about 90% of the earnings penalty among graduates occurs between personal interviews and job offers (Figure 1).

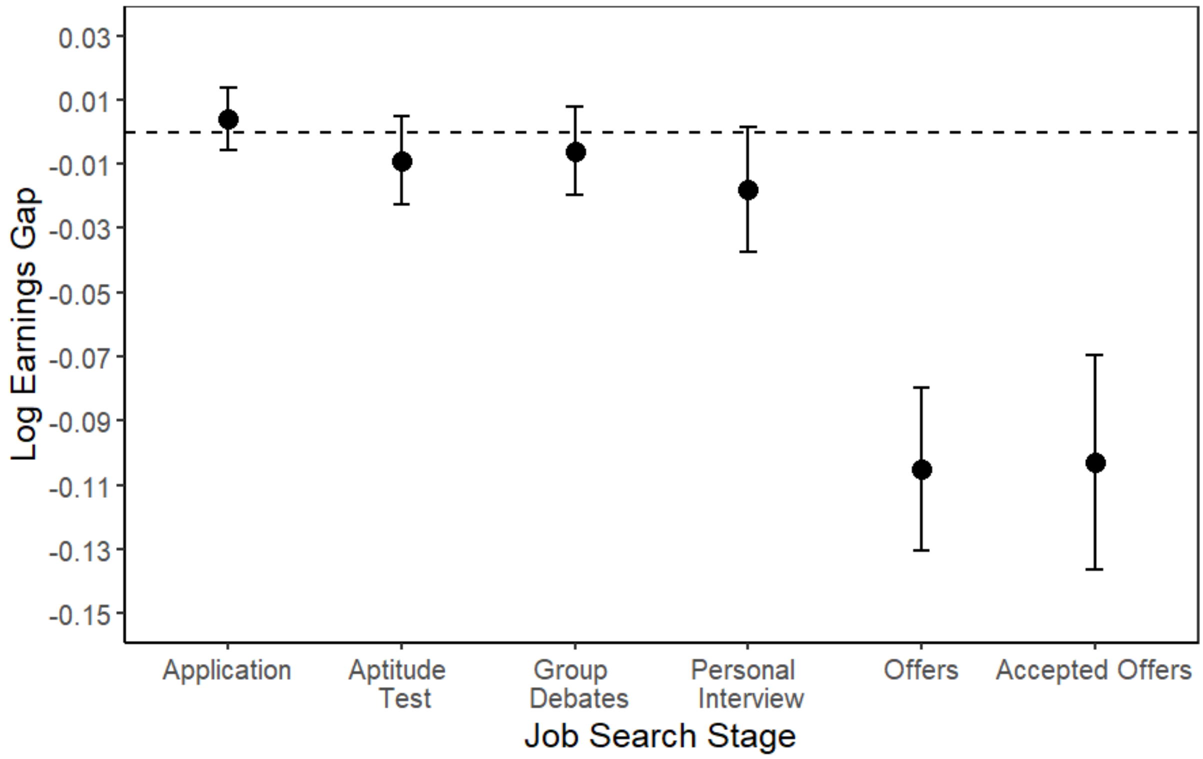

Figure 1: The Caste Earnings Penalty

Notes: Figure 1 shows the log earnings gap across castes at successive stages of job search. Each coefficient in the figure is represented by a black dot and reports the percentage difference in the average salary at each job search stage between disadvantaged and advantaged castes. The vertical bars are 95% confidence intervals. These regressions include controls. Note that firms post uniform, job-specific (not match-specific) wages that are non-negotiable over the course of job search. Thus, the caste earnings penalty among graduates can be decomposed by comparing the caste difference in the composition of jobs that remain in contention at each stage of job search.

7.3 Almost the Entire Caste Penalty in the Success Rate Occurs At Personal Interviews

Throughout the paper, "success rate" is defined as the probability that a given student i makes it past screening round s for job j in year t.

The role of personal interviews in explaining caste disparities in hiring is much stronger if, instead, one were to quantify the relative contribution of each of the four screening rounds in explaining the lower success rate of disadvantaged castes compared to advantaged castes at the job offer stage. Disadvantaged castes are about 30 percentage points (or about 115%) less likely to get job offers than advantaged castes. More than 95% of this caste penalty in the success rate can be attributed to the personal interview round (Table 6). Therefore, disadvantaged castes are almost as likely to succeed as advantaged castes at any of the three screening rounds prior to personal interviews.

Since jobs post uniform, job-specific salaries, the results reported in Section 7.2 do not necessarily account for heterogeneity across jobs. To show caste differences in the success rate at each screening round, I run standard linear probability regressions with job fixed effects. The results are reported in Table 6, where an observation in each of the columns of Table 6 is a student who remains in contention at a given screening round in a given job in a given year. For example, all students in the first column of Table 6 are applicants. Similarly, all students in the second column of Table 6 are written test takers and so on. Specifically, I estimate the following linear probability model:

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

where $i \in \mathcal{I}$ indexes students, $j \in \mathcal{J}$ indexes jobs, Search Stage $\in$ {Application, Aptitude Tests, Group Debates (GD), Personal Interviews, Offers, Accepted Offers}, γ_t denotes year fixed effects, and k_j denotes job fixed effects.[23] The controls are standard characteristics such as college GPA, college entrance exam score, previous labor market experience, major, and degree. Table 6 shows that the caste difference in the success rate is negligible at any of the three screening rounds prior to personal interviews. At the job offer stage, however, there is a substantial caste penalty. Disadvantaged castes are about 30 percentage points less likely to get job offers following personal interviews, a nearly 115% lower probability of success relative to the mean success rate of 26% at interviews.

As mentioned in Section 5.3.2, it is typical for a firm to conduct a common group debate or personal interview for its multiple job positions and these jobs are counted separately in the regressions reported in Table 6. For example, a student who participates in a common group debate conducted by Google for each of its five advertised job positions will be considered to have participated in five separate group debates, one for each job, as he could potentially receive five different job offers. Similarly, a student who participates in a common personal interview conducted by Uber for each of its two advertised job positions will be considered to have participated in two separate interviews.

In summary, Sections 7.2 and 7.3 show that caste disparities in hiring outcomes in the multi-stage job search process primarily arise after the final screening round comprising personal interviews that often assess fit. These results suggest that policies informing applicants about job opportunities, modifying student preferences, or improving performance at university are unlikely to meaningfully mitigate caste disparities. Other policy implications of these findings, in India and elsewhere, are discussed in Section 11.

[23] See Section 5.2 for how a "job" is defined.

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

	Success App (1)	Success Test (2)	Success Debate (3)	Success Interview (4)
Disadv. Caste	-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

Notes: Table 6 shows the caste difference in the success rate at each stage of job search. Each observation is a student who remains in contention at a given screening round in a given job in a given year. Each column reports a separate linear probability regression and includes controls. Note that some firms conduct the same screening round (written tests, group debates, interviews) for multiple job positions. However, since these job positions within the same firm represent different jobs, the regressions consider these jobs and the corresponding advance-or-out decisions made by them separately. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

7.4 Caste Disparities Are Concentrated Following Personal Interviews in About 50% of Jobs

Estimating a slightly different version of Equation 2 separately for each job, I find two results (not shown) in about 50% of the jobs in my sample. First, caste disparities following personal interviews in these jobs are at least as large as the average caste penalty in the success rate at the job offer stage reported in Table 6. Second, there are negligible caste disparities in the success rates at any of the three screening rounds prior to personal interviews. Therefore, caste disparities are concentrated following personal interviews in about 50% of the jobs in my sample.

7.5 Modest Heterogeneity Across Job Types in the Earnings Gap And the Success Rate At Personal Interviews

There is little heterogeneity by job type in the caste earnings gap. Recall that the average caste earnings gap among graduates is about 10%. The earnings gap is -0.107 (0.027) log points among client-facing jobs that comprise about 25% of available jobs. The gap is -0.089 (0.022) log points among non-client-facing jobs that comprise about 75% of available jobs (Appendix Table OA.1).

There is also little heterogeneity by job type in the caste penalty in the success rate at any screening round. There is a negligible caste penalty in the success rate at any of the three screening rounds prior to personal interviews among both client-facing and non-client-facing jobs. However, the caste penalty in the success rate at the personal interview round is about 30 percentage points among both client-facing and non-client-facing jobs. The mean success rate at personal interviews is also about 26% among both client-facing and non-client-facing jobs (Online Appendix Tables OA.4 and OA.5). Therefore, disadvantaged castes have about a 115% lower probability of getting a job offer after personal interviews regardless of job type.

7.6 Caste Differences in Job Tenure and Promotion Rates

Caste disparities in initial hiring outcomes are not aligned with productivity on the job, as disadvantaged castes often have better job performance once hired. While I do not observe direct measures of personal interview performance, I tackle this limitation by collecting other, post-hiring measures of job performance in the same jobs that students received and accepted through the elite college's regular placement fair for full-time jobs. Linking the sample to data from a major employment-focused platform provides two measures of job performance: tenure and promotion rates. I can link over 98% of the sample of students who received jobs through the job fair. Job tenure is measured in months. Promotion is defined as the probability a student who received a job through the on-campus job fair is promoted (further defined in Section 5.5). To compare caste differences in job performance, I run different specifications of the following linear regression:

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

where $i \in \mathcal{I}$ indexes students, $j \in \mathcal{J}$ indexes jobs, γ_t denotes year fixed effects, and k_j denotes job fixed effects.[24] The controls are standard characteristics such as college GPA, college entrance exam score, previous labor market experience, major, and degree. Recall that the primary focus of the paper is to study whether personal interviews are avenues where employers (unwarrantedly) put a higher threshold on disadvantaged castes through subjective assessments of "fit." Therefore, as in Section 7.1, I focus the discussion of caste differences in tenure and promotion rate only among the sample of students who secured employment through the college's regular placement fair for full-time jobs, and therefore underwent a standardized hiring process comprising four screening rounds per job (see Section 7.1 for more details on how the sample is divided between "job getters" and "non-job getters"). I discuss selection and its impact on the results of the outcome tests in Section 7.6.1.

On average, disadvantaged castes have 0.145 (0.023) or about 15% higher job tenure than advantaged castes (Column (1) of Table 7). The average tenure duration in the first job is about 2 years. While there is precedence of using job tenure as a measure of job productivity (Hoffman et al., 2018), it could also present issues. For example, disadvantaged castes may have higher tenure because they have poorer outside options.

To improve upon my measure of job performance, I compare the caste difference in promotion rates among the jobs received by students through the college's job fair. Promotions are a particularly good measure in this setting because entry-level jobs at elite multinationals typically have rigid salary structures, in India and elsewhere (Hazell et al., 2021; Sarsons et al., 2022). I show that, conditional on job tenure and other controls, disadvantaged castes are 0.096 (0.026) percentage points more likely to be promoted in a given job (Column (5) of Table 7). Relative to the average promotion rate of about 70%, disadvantaged castes are about 15% more likely to be promoted than advantaged castes. Similar results hold when the promotion regression includes either only job fixed effects or only job and year fixed effects (Columns (2)-(3) of Table 7).

I do not estimate promotion rates for the marginal candidate. To deal with potential issues due to inframarginality, I construct a measure of job selectivity. "Job selectivity" is proxied by calculating the cumulative proportion of job applicants from the elite college cut by a given job by the end of the final screening round comprising informal personal interviews. For example, if a job makes job offers to only 5% of its initial applicant pool, the cumulative proportion of applicants cut by the job by the end of the personal interview round is 95%. Using this measure, I categorize jobs according to their selectivity and quantify caste differentials in promotion rates for jobs in the top 75%, 50%, 25%, and 10% of selectivity.

I show that disadvantaged castes have about a 15% higher promotion rate, *even among the most selective*

[24] See Section 5.2 for how a "job" is defined.

jobs. Columns (6)-(9) of Table 7 show that the coefficients in all specifications are around 15% in magnitude. Except for jobs in the top 10% of selectivity, where the coefficient on disadvantaged caste is 0.084 (0.109), these coefficients are also statistically significant. Note that the result for jobs in the top 10% of selectivity is almost identical to that in the full sample, but the standard error is up based on the decline in sample size.[25] Disadvantaged castes also have higher promotion rates among both client-facing and non-client-facing jobs (Columns (10)-(11) of Table 7). Overall, the results in Table 7 suggest that caste disparities in my setting are consistent with caste discrimination in hiring under reasonable assumptions about the inframarginal distribution of talent.

Admittedly, outcome tests are not a panacea and often come with a host of known issues, some less severe than others. Omitted variable bias typically does not significantly affect outcome tests. A decision-maker should have already considered all factors that influence "success." Therefore, analyzing success itself should already control for those factors (Becker, 1957). On the other hand, inframarginality is a well-known problem that could affect the results of outcome tests (Ayers, 2002; Ayres, 2010; Hull, 2021; Mogstad et al., 2023). However, my results showing that disadvantaged castes often outperform advantaged castes even among the most selective jobs should ease some concerns related to inframarginality. Furthermore, the outcome tests results should be seen as supporting the broader story of caste disparities in elite jobs, and therefore should be viewed jointly with the evidence documented in previous sections, rather than as a be-all and end-all exercise.

[25] A standard equality of coefficients test can be used to compare coefficients across the specifications reported in Table 7.

Table 7: Caste Differences in Job Tenure and Promotion Rate

	Dependent variable:										
	log(Tenure Duration)	Promoted	Promoted	Promoted	Promoted	Promoted at Top 75%	Promoted at Top 50%	Promoted at Top 25%	Promoted at Top 10%	Promoted at Client Facing	Promoted at Non-Client Facing
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Disadv. Caste	0.145***	0.140***	0.137***	0.129***	0.096***	0.104***	0.112***	0.130**	0.084	0.107**	0.082**
	(0.023)	(0.024)	(0.025)	(0.026)	(0.026)	(0.029)	(0.038)	(0.062)	(0.109)	(0.044)	(0.032)
GPA	0.077***			0.100***	0.064**	0.055	0.053	0.054	0.027	0.077	0.056
	(0.027)			(0.030)	(0.031)	(0.035)	(0.051)	(0.082)	(0.206)	(0.057)	(0.037)
Tenure Duration					0.084***	0.082***	0.076***	0.065**	0.115**	0.065***	0.097***
					(0.012)	(0.014)	(0.019)	(0.027)	(0.049)	(0.021)	(0.015)
Job FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓
Other Controls	✓			✓	✓	✓	✓	✓	✓	✓	✓
Mean Y	25.77	0.684	0.684	0.684	0.684	0.667	0.668	0.711	0.73	0.687	0.682
Observations	2,489	2,489	2,489	2,489	2,489	2,012	1,246	623	252	835	1,654
Adjusted R^2	0.119	0.183	0.183	0.186	0.205	0.203	0.216	0.185	0.127	0.194	0.214

Notes: Table 7 shows caste differences in job performance measures among students who graduated with full-time jobs through the college's on-campus job fair. All of these students had to succeed in the final round of subjective, non-technical personal interviews to obtain jobs. Column (1) reports caste differences in tenure. This regression includes both job and year fixed effects as well as controls such as college GPA, degree, and major. Job tenure is measured in months. Column (2) reports caste differences in promotion rates only with job fixed effects as controls. Column (3) reports caste differences in promotion rates only with job and year fixed effects as controls. Column (4) reports caste differences in promotion rates with job and year fixed effects as well as controls such as GPA, degree, and major. Column (5) reports the same regression as Column (4) but with job tenure (standardized) as an additional control. Columns (6)-(9) show caste differences in promotion rates with controls among jobs at different quantiles of "job selectivity," defined as the cumulative proportion of applicants cut by a given job by the end of final round personal interviews. Columns (10)-(11) show caste differences in promotion rates with controls among client-facing and non-client-facing jobs, respectively. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

7.6.1 Caste Differences in Job Tenure and Promotion Rates Are Conservative

The results in Table 7 could be affected by sample selection. The first source of sample selection arises from the absence of job tenure and promotion data for some students who received full-time jobs, leading to unmatched data. However, among the sample of students who got full-time jobs by participating in the college's job fair, I am not able to match outcomes data for less than 2% of the sample.[26] Further, relative to this sample of initial full-time job getters, those not in the tenure and promotion regressions are much more negatively selected among advantaged castes (not shown). In other words, among those included in the promotion and tenure regressions, advantaged castes are relatively more positively selected and thus the results reported in Table 7 are likely underestimates.

There is a yet more severe source of sample selection that could affect the results in Table 7. About 45% of those in the regressions reported in Table 7 leave the first job they received and accepted through the college's job fair at some point during the timeframe for which promotion and tenure data is collected.[27] Students might leave the first job for a variety of reasons, including changing career tracks; pursuing an MBA, a career in the elite civil services, or an advanced degree; or lack of opportunities in the current job. Sample selection among "first job leavers" could cause caste differences in both job tenure and promotion rates to be overestimated. Specifically, if the "best" advantaged castes leave the first job, then the comparisons in Table 7 could be between disadvantaged castes and relatively more negatively selected advantaged castes, resulting in caste differences in job promotion rates to be overestimated. However, Online Appendix Table OA.6 shows that relative to the sample of students who stay in the first job, those who leave are much more negatively selected among advantaged castes. In other words, among those who stay in the first job received and accepted through the college's job fair, advantaged castes are relatively more positively selected. Therefore, caste differences in both job tenure and promotion rates are likely underestimated.

7.6.2 Why Do Disadvantaged Castes Have Better Job Performance?

It is beyond the scope of this paper to try to formally understand why disadvantaged castes have better job performance. However, I briefly discuss possible reasons for my finding.

Disadvantaged castes who qualify for selective elite entry-level jobs might have to face greater hurdles, and therefore might be of higher quality. Consistent with my findings, Anzia and Berry (2011) and Ferreira and Gyourko (2014) show that "equally qualified" female politicians in the U.S. outperform their male counterparts. Specifically, Anzia and Berry (2011) find that congresswomen secure about 10% more spending from federal discretionary programs than congressmen. Ferreira and Gyourko (2014) also find that female mayors have about a 15% higher "incumbent effect" (chance of winning re-election) than comparable male mayors. Relatedly, Hengel and Moon (2023) show that at leading economics journals, female-authored papers are held to higher standards and receive about 20% more citations than comparable male-authored papers. These performance differences are comparable to those reported in Table 7.

In the Indian context, and thus closest to my own setting, Bhavnani and Lee (2021) show that disadvantaged castes who qualify for elite bureaucratic positions in the Indian Administrative Service (IAS) *with*

[26] The size of this sample is 2,527 - 2,489 = 38 students, which is only about 1.5% of the sample that graduated with full-time jobs obtained through the college's job fair.

[27] About 50% of advantaged castes and 40% of disadvantaged castes leave the first job at some point during the timeframe for which promotion and tenure data is collected.

affirmative action perform similarly to advantaged castes.[28] Further, the authors show that a substantial portion of disadvantaged castes outperform comparable advantaged castes. Specifically, the authors show that disadvantaged castes who qualify for the IAS *without* affirmative action have about 20% better performance than advantaged castes.[29] This performance difference is also comparable to that reported in Table 7.

Recruitment of IAS officers typically occurs in two stages. The first stage is a written component that is followed by a subjective, in-person oral interview, or "personality test," that is conducted in English by an interview board comprising primarily advantaged castes. Bhavnani and Lee (2021) document that disadvantaged caste IAS officers have significantly worse personal interview scores, even after controlling for written test performance. They also show that total exam score (of which personal interview score is a big component) and IAS officer performance are weakly correlated. Moreover, conditional on exam scores, disadvantaged caste IAS officers recruited without affirmative action perform better than advantaged castes. Bhavnani and Lee (2021) suggest that not only could personal interviews provide better information about caste status (the written test is anonymous) but also that scores obtained through such subjective evaluations might be biased. Thus, personal interviews might understate "true" candidate quality.[30]

Similarly, in my setting, it is plausible that (conditional on reaching the personal interview round) personal interview scores are weakly correlated with productivity. Recall, nearly 75% of the jobs in my setting are non-client-facing. In such jobs, it is plausible that job performance is likely evaluated on more objective criteria such as code quality than subjective aspects of "fit." Therefore, it is plausible that scores on soft aspects of cultural fit could understate candidate quality, especially among the already qualified pool of students who make it to the personal interview round of highly selective jobs.

7.6.3 Interviews Are Common Avenues Where Disparities Emerge

In this section, I briefly mention studies from sociology, psychology, and organizational decision making that have examined topics like the one in this paper. I conclude the section by clarifying my contribution over these studies.

Many studies in sociology, psychology, organizational decision making and even medicine have shown that unstructured interviews, evaluations, or other forms of subjective assessments are common avenues where disparities emerge (see e.g., Meehl, 1954; Carroll et al., 1982; DeVaul et al., 1987; Highhouse, 2008; Rivera, 2012; Dana et al., 2013; Rivera, 2015). For example, Meehl (1954) conducted a meta-analysis of studies comparing human and statistical predictions, and showed that statistical prediction was almost always superior. Carroll et al. (1982) showed that parole boards' subjective judgments about the risk of recidivism did not produce accurate predictions of post-release outcomes. DeVaul et al. (1987) showed that medical students who were initially rejected based on subjective evaluations (e.g., interviewer ratings and committee ratings) were just as successful as their non-rejected peers when they were later admitted to

[28] In the IAS, disadvantaged castes can qualify for bureaucratic positions either with or without affirmative action. In Bhavnani and Lee (2021), "performance" is measured by calculating the number of households that received at least 100 days of guaranteed employment under a major anti-poverty program of the Indian government.

[29] Specifically, Bhavnani and Lee (2021) show that disadvantaged castes who qualify for the IAS without affirmative action improve the number of households that received at least 100 days of guaranteed employment under a major anti-poverty program of the Indian government by 0.1 standard deviations. This effect size is about 20% relative to the mean and statistically significant at the 10% level.

[30] The authors state that "[t]he fact that disadvantaged group IAS recruits perform poorly on the interview portion of the recruitment exam, where it is relatively easy to guess caste identity, rather than the more objective written portions of the exam, points to the specific stage at which candidate quality is understated" (Bhavnani and Lee, 2021, p. 4-5).

medical school because a new law required the medical school to increase its class size late in the recruiting season. Highhouse (2008) argued that evaluators are stubbornly reliant on "intuition and subjectivity in employee selection" and that they have a bias against decision aids such as written tests. He further stated that, "[a]lthough it is commonly accepted that some (employment) interviewers are better than others, research on variance in interviewer validity suggests that differences are due entirely to sampling error." Rivera (2012) interviewed hiring professionals in elite U.S. firms and found that evaluators openly admitted that "cultural fit" was the most important criterion for them in their interview evaluations. Such requirements were commonly expressed in elite law firms (75%), investment banks (65%), and consulting firms (40%). Unstructured interviews could even make other predictors of applicant success less useful by adding noise. For example, Dana et al. (2013) show that prior GPAs predicted future GPAs very well with a correlation of 0.65, but when unstructured interview-based evaluations were added in, the correlation reduced to 0.31. To summarize, unstructured interviews, evaluations, or other forms of subjective assessments are common avenues where disparities emerge. Therefore, the goal of this paper is not to posit unstructured interviews as a "new" mechanism for disparities.

This paper goes beyond showing that personal interviews are common avenues where disparities emerge, quantifying the real effects of disparities due to such subjective screening practices in earnings and access to elite jobs. The lagging representation of disadvantaged castes in elite jobs, despite broad-based affirmative action in university admissions in India, could be partly driven by screening processes which allow employers to discriminate against disadvantaged castes. Even at American workplaces, screening practices have limited the economic mobility of disadvantaged castes, often by allowing managers to covertly retaliate against employees (California v. Cisco Systems, Inc., 2022). Yet, little is known about how quantitatively important screening processes are in driving caste disparities in access to the "elite," or what aspects of candidate screening drive these disparities. Adding to these challenges is the fact that recruitment practices of elite jobs are often non-transparent (Rivera, 2015). My paper addresses these issues by showing that the final screening round consisting of subjective, non-technical personal interviews that often assess "fit" could create barriers for disadvantaged castes in accessing elite jobs. Specifically, I show that nearly 90% of the caste disparities in hiring arise due to personal interviews (Figure 1 and Table 6). I also argue that this screening round is where employers are most likely to learn about caste. Further, information learned during personal interviews is weakly correlated with job performance. Instead, personal interviews, a common component of elite job recruitment in India and elsewhere, are a large contributor to unwarranted socioeconomic (caste) disparities in the labor market. Such disparities could be motivated by taste-based, statistical, customer, or coworker discrimination, or other factors. To my knowledge, this paper is the first to quantify *how much* of the hiring disparities, however motivated, among elite jobs in a real world and high-stakes setting, could be attributed to subjective assessments of fit done through personal interviews.

8 Caste Penalty and Caste Revelation

The results documented in Section 7 beg an obvious question: when does caste get revealed to employers? In this section, I argue that the caste penalty in hiring outcomes is closely linked to caste revelation.

8.1 What Does the Employer See At Each Screening Round?

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

1. **Application Reading.** No part of the application process requires students to declare their caste status or their incoming college entrance exam scores. At this round, the employer has access to CVs that are written in a standardized format prescribed by the placement office and typically comprise information about student names, college GPA, work experience, and extracurricular activities.

2. **Written Tests.** Written tests about technical skills related to the job are conducted online. Therefore, the only additional information employers obtain, relative to the previous application screening round, is the written test scores.

3. **Group Debates.** At this round, employers finally get to observe—but do not extensively interact with—students after an initial setup. They observe additional information like facial features, skin tones, accents, dialects, and demeanor.

4. **Personal interviews.** These interviews represent the only extensive one-on-one interaction between the employer and the candidate. Conversations during these interviews are typically free-form or unstructured. Such interviews are avenues where informal conversations about shared experiences, personal hobbies, educational qualifications of family members, father's job, neighborhood of residence, preference for living in a cosmopolitan city, desire for traveling, and other aspects of "cultural fit" are considered par for the course (Section 6).

As mentioned in Section 5.6, I do not have access to student CVs or performance measures at any screening round. I address potential limitations due to lack of access to such data and their implications for my findings in Sections 9 and 10.

8.2 Quantitative Evidence

I provide quantitative evidence linking the caste penalty to caste revelation. In a small subsample where last names provide clearer caste signals, I show that the penalty emerges much earlier—during the application reading round. However, the *overall* caste penalty during application reading is negligible because most of the sample (about 85%) has non-distinctive caste names.[31] Instead, the caste penalty primarily emerges during personal interviews, where students are screened based on background, hobbies, and cultural fit: characteristics strongly correlated with caste.

I show this result through Figures 2 and 3. Figure 2 shows a different version of Figure 1 by splitting the sample into two groups: students with clearly identifiable caste names and those without. I separately estimate Equation 1 for these two samples and report the coefficients with 95% confidence intervals in Figure 2. The right figure in the top panel of Figure 2 clearly shows that a substantial caste earnings penalty emerges immediately during the application reading round in the sample with identifiable caste names. In contrast, the left figure in the top panel of Figure 2 shows that no meaningful caste earnings penalty emerges during the application reading round in the sample with nonidentifiable caste names (compare the coefficients on

[31] Recall that while many surnames are indicative of caste, most of them do not convey caste signals, especially in urban-educated settings due to migration and geographical variation in naming conventions. Also see Footnote 3.

the top of "Application" and "Aptitude Test" in both figures in the top panel of Figure 2). Online Appendix Tables OA.7 and OA.8 report the regressions that produce the coefficients shown in Figure 2.[32]

Further, I compare success rates at the application reading round for students with identifiable and non-identifable caste names. To do so, I first extend the definition of "job selectivity" introduced in Section 7.6 to include the cumulative proportion of students cut by a given job by the end of *any* screening round. Next, I show that in the sample of students with identifiable caste names, disadvantaged castes have lower success rates than advantaged castes at the application reading round, essentially across the *entire distribution* of job selectivity. However, no such caste penalty in the success rate exists in the sample of students with nonidentifiable caste names. I show these results through Figure 3. Specifically, for each job $j \in \mathcal{J}$, I estimate the following linear probability model:

$$\text{Pr}(\text{Success}_{ijt}^{\text{Job Search Stage}}) = \alpha_j + \beta_j \times \text{Disadv. Caste}_i + \text{Controls}_i + \gamma_t + \text{error}. \tag{4}$$

where $i \in \mathcal{I}$ indexes students and γ_t denotes year fixed effects. The controls are standard characteristics such as college GPA, college entrance exam score, previous labor market experience, major, and degree. In Figure 3, the x-axis shows the cumulative cuts made by jobs by the end of the application reading round.[33] The y-axis shows the conditional (i.e., with controls) caste gap in the success rate at the application reading round. A caste gap in the success rate below zero means that disadvantaged castes have a lower success rate or a higher cut rate than advantaged castes ("caste disadvantage"). The top panel shows that, in the full sample of students, there are negligible caste differences in the success rate at the application reading round across jobs with varying levels of selectivity. The average caste penalty in the success rate at the application reading round in the full sample (top panel, thick black line) is similar to the one reported in Column (1) of Table 6. The bottom panel shows that clearly identifiable disadvantaged castes have a substantially lower success rate than comparable advantaged castes at the application reading round, essentially across the *entire distribution* of job selectivity. The caste penalty in the success rate at the application reading round in the sample with identifiable caste names is about 15 percentage points or about 30% relative to the mean success rate at the application reading round for this sample. In contrast, no notable caste penalty in the success rate at the application reading round exists in the sample of students with nonidentifiable caste names.

[32]While not significant, the non-monotonicity of the coefficients with respect to the cognitive testing stage in the right figure in the top panel of Figure 2 is consistent with objective evaluation methods (e.g., written tests) potentially mitigating disparate outcomes resulting from relatively more subjective evaluation methods (e.g., application reading).

[33]Note that since the application reading round is the first screening round, the "conditional cuts" are identical to the "cumulative cuts." For more details, see Section 5.3.1.

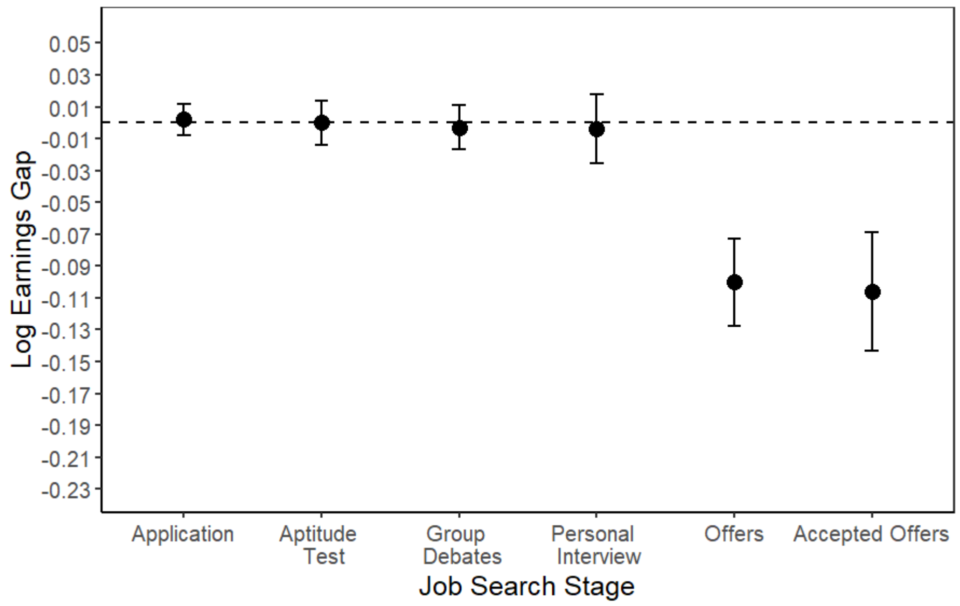

(a) Earnings Gap across Castes at Each Job Search Stage in the Nonidentifiable Sample

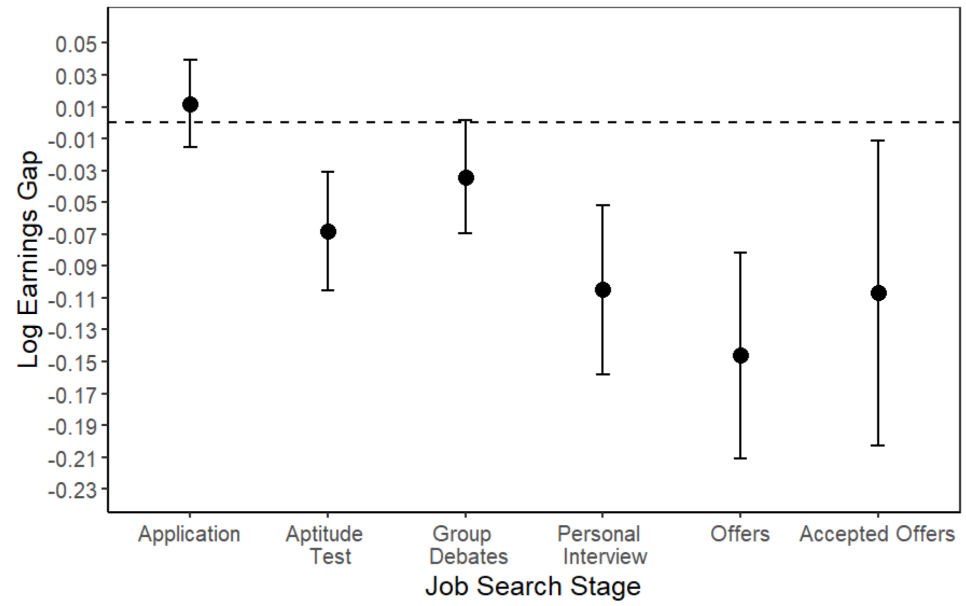

(b) Earnings Gap across Castes at Each Job Search Stage in the Identifiable Sample

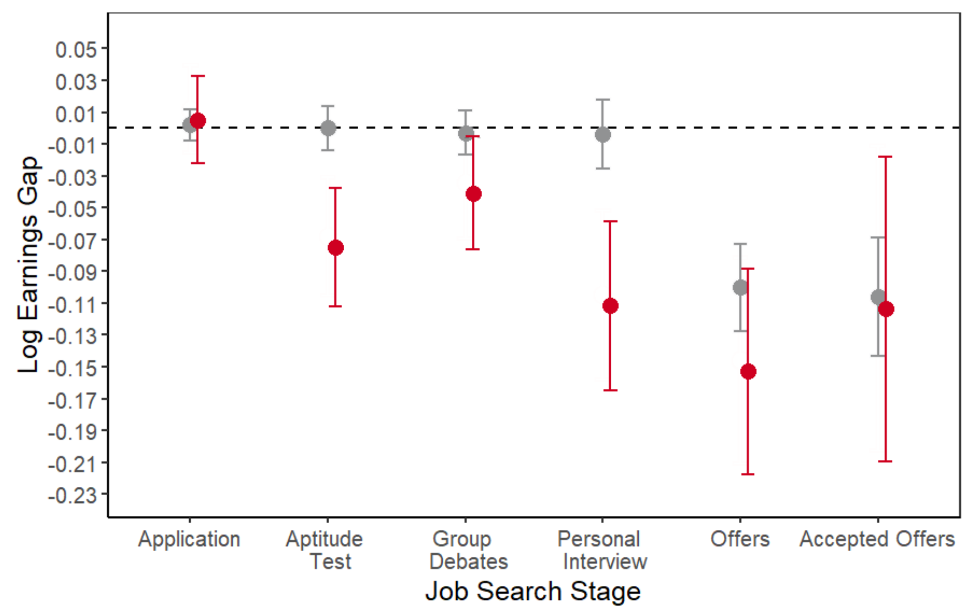

(c) Earnings Gap across Castes at Each Job Search Stage in the Identifiable (Red) and Nonidentifiable (Gray) Sample

Figure 2: Caste Earnings Penalty in the Identifiable and the Nonidentifiable Sample

Notes: The top panel of Figure 2 shows the log earnings gap across castes at successive stages of job search separately for both the identifiable (right figure) and the nonidentifiable sample (left figure). The same results are shown together for the identifiable sample (red) and the nonidentifiable sample (gray) in the bottom panel of Figure 2. Each coefficient in the figure is represented by a black dot and reports the percentage difference in the average salary at each job search stage between advantaged and disadvantaged castes. The vertical bars are 95% confidence intervals. These regressions include controls. Note that firms post uniform, job-specific (not match-specific) wages that are non-negotiable over the course of job search. Thus, the caste earnings penalty among graduates can be decomposed by comparing the caste difference in the composition of jobs that remain in contention at each stage of job search.

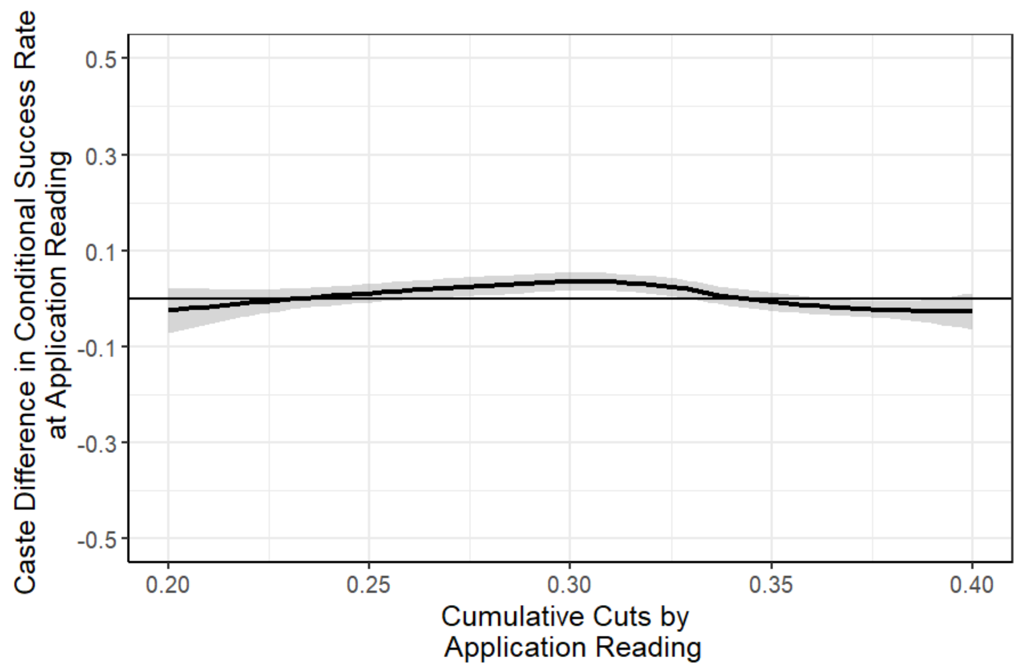

(a) Caste Gap in Conditional Success Rates versus Cumulative Cuts at Application Reading in the Full Sample

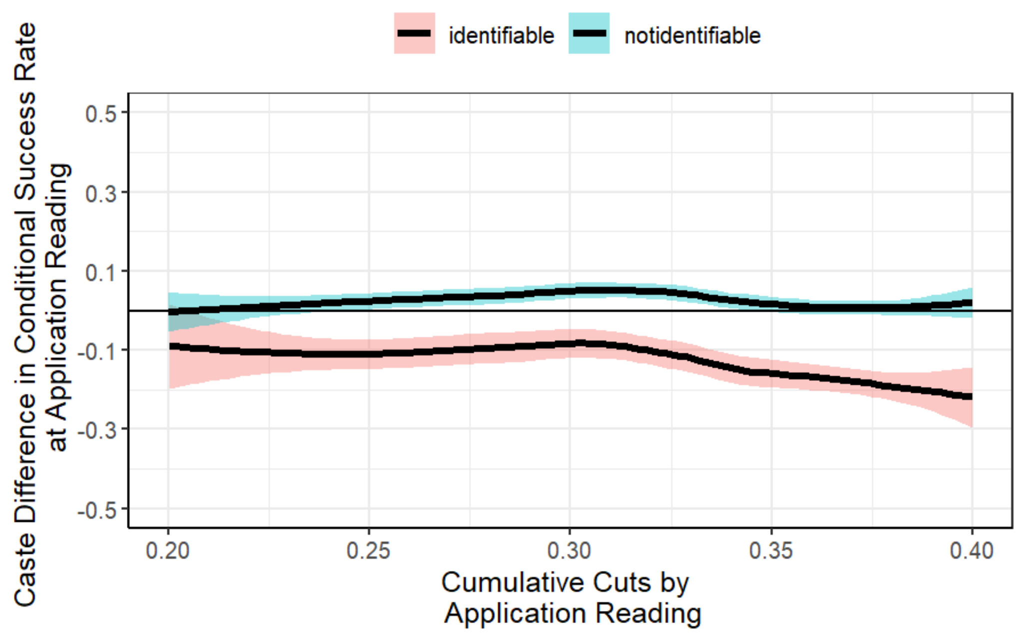

(b) Caste Gap in Conditional Success Rates versus Cumulative Cuts at Application Reading in the Identifiable and Nonidentifiable Sample

Figure 3: Caste Penalty in the Conditional Success Rate At the Application Reading Round in the Identifiable and the Nonidentifiable Sample

Notes: Figure 3 shows the caste difference in the conditional success rate against the cumulative cuts made by jobs by the end of the application reading round. Smoothed conditional means (thick black lines) are calculated using the "loess" function in R. The shaded areas in gray, blue, and red denote 95% confidence intervals. The thin solid horizontal line denotes a y-intercept of 0. Panel (a) is constructed from the full sample. Panel (b) is constructed separately for the sample with identifiable (red) and nonidentifiable (blue) caste names. The x-axis shows the cumulative proportion of students cut by jobs by the end of the application screening round. The y-values are calculated by estimating a standard linear probability model with controls for each job and then plotting the coefficient on disadvantaged caste from these regressions on the y-axis.

8.3 Qualitative Evidence

The quantitative evidence presented in Section 8.2 linking the caste penalty to caste revelation is also consistent with the growing role of socioeconomic cues or "soft" factors such as cultural fit in aiding caste identification in elite, urban-educated settings. While admittedly weaker than direct data-based evidence, it is still important to qualitatively discuss the nuances of caste identification and how it typically operates in modern, urban-educated India. Further, the qualitative evidence should align with the quantitative evidence presented in Section 8.2 above linking the caste penalty to caste revelation.

Caste is unlikely to be signaled in the full sample through names during application reading or online tests given large regional variation in names and naming conventions, the elite college drawing students from myriad regions, and hundreds of applicants per job. For example, surnames like "Singh," "Sinha," "Verma," "Chaudhary," "Mishra," and "Das" are shared across castes (Anthropological Survey of India, 2009). Moreover, in a recent audit study based on firms in the New Delhi area, Bertrand et al. (2009) state that the "enormous regional variations [in last names] mean that the precise coding of a particular last name is unlikely to be familiar to people from a different linguistic region of India." Furthermore, naming conventions also differ significantly across regions. For example, many South Indians do not have conventional last names. Rather, for them, personal (first) names often perform the roles of traditional "surnames" (Jayaraman, 2005). These explanations are consistent with my finding that only about 15% of the student sample in this elite college has clearly identifiable caste names. Similarly, it is also unlikely that caste status is known during written aptitude tests or "technicals," as these tests are typically conducted online. Other than scores obtained during written tests, employers do not observe any additional information about candidates relative to the previous application screening round.[34]

Caste identification is also unlikely to be reliable during group debates. Recall, these debates are conducted among large groups of 30-35 total students who either argue for or against a given topic. In the data, students at group debates are about evenly split between castes. At this round, employers finally get to observe but do not extensively interact with students after an initial setup. They observe additional information like facial features, skin tones, accents, dialects, and demeanor. However, scholars have argued that there is no association between skin color and caste, especially since Indian skin color is influenced mostly by geographic location rather than caste status (Mishra, 2015; Parameswaran and Cardoza, 2015). Among educated elites, like those in my sample, English has emerged as a caste-neutral language with little to no remnants of caste dialects that are prevalent in most Indian languages (Kothari, 2013; Ransubhe, 2018). Rather, perception of accent variation among young, English-speaking university graduates in India is linked to broad regional factors, such as a North Indian versus South Indian accent, instead of caste status (Wiltshire, 2020).[35] Relatedly, multiple studies show substantial "convergence" in English accents and dialects in broad, urban-educated Indian regions due to migration, urbanization, and improved levels of education (Wiltshire, 2020). To summarize the discussion so far, cues such as last names, accents, dialects, and skin color are likely to be weak signals of caste status, especially in settings such as elite colleges characterized by a lot of "mixing," as they draw students from myriad geographical regions.

[34] Even assuming that employers actually know the distribution of college GPA by college degree and year, college GPA is a noisy signal of caste status outside of the tails of the distribution (Online Appendix Figure OA.1).

[35] The overwhelming influence of regionality in common parlance is perhaps most clearly expressed by Gumperz (1961), who states that a "high-caste villager may speak the same form of urban Hindi as his untouchable neighbor."

However, conversations during non-technical personal interviews likely provide additional information that could reveal caste. These interviews primarily screen on parental background, neighborhood, and other aspects of "cultural fit." Such background characteristics are strongly correlated with caste in elite, urban-educated India (see Section 6; Deshpande, 2011; Bhavnani and Lee, 2021; California v. Cisco Systems, Inc., 2022). These caste-correlated characteristics are generally hard to fake because caste networks are deeply entrenched (Beteille, 1965, 1969; Jodkha, 2017). Recent experimental studies among college educated students in urban India have also shown that background characteristics convey almost perfect signals of caste status relative to surface-level attributes (e.g., last names, facial features, and dialects) that are highly evenly shared across castes, and therefore, obfuscate caste identification (Mamidi, 2011).

In light of the above discussion, the descriptive evidence presented in Section 7 advances the literature on the detection and measurement of labor market disparities. As intimated above, disparities based on socioeconomic cues or "soft" factors like cultural fit are likely to become more salient in elite, urban-educated settings because standard characteristics by which to differentiate groups become less perceptible as these settings become more multi-ethnic and diverse (Loury, 2002; Freeman et al., 2011; Gaddis, 2017). Therefore, collecting data on various stages of job search, each having the potential to reveal different socioeconomic cues, is crucial to better detect group disparities in the context of elite hiring. I provide an important example of the kinds of data future researchers may have to collect in such settings to better detect disparities from outwardly neutral screening practices. Even in much of the Western world, as racial boundaries blur, "caste" discrimination is likely to become a bigger part of the discourse globally (Austen-Smith and Fryer, 2005; Eguia, 2017; Kim and Loury, 2019; Rose, 2022). Indeed, as stated by Rose (2022), group differences could often "emerge *endogenously* as an outcome of stereotypes that, in equilibrium, engender real disparities across groups" and therefore reinforce perceptions of racial, class, or caste differences. Discriminatory practices in elite jobs can often be coded through hiring managers valuing subjective aspects of fit, a practice that often favors socioeconomically advantaged groups. Yet, recent work studying the role of hiring discretion has either exclusively focused on only low-skilled jobs or provided correlational evidence between callback disparities and HR practices being more subjective (Hoffman et al., 2018; Kline et al., 2022).

My empirical approach to capture the caste penalty through a reduced-form caste coefficient that potentially captures various sources of caste disparities also helps advance recent research that argues for a constructivist understanding of group identities, instead of treating them as immutable facts (Hull et al., 2022; Rose, 2022). Such an approach is crucial to better understand "caste," classifications of which are rooted in the economic, political, and material history of India (Beteille, 1965, 1969). In addition, as argued in this section, perceptions of caste in elite, urban-educated India are guided by a myriad of socioeconomic cues. Such perceptions parallel those of social class in other contexts, especially Britain and the United States.

8.4 Preferred Interpretation of the Caste Penalty

Given the evidence documented in Sections 8.2 and 8.3, my preferred interpretation of the caste penalty is that it stems from employers directly penalizing caste. Employers may penalize caste due to taste-based, statistical, customer, or coworker discrimination, unintentional biases due to homophily, or other factors. I do not distinguish between these sources of discrimination. Alternatively, one could also rationalize caste disparities as stemming from a scenario in which employers do not directly penalize caste. Under this scenario, employers could value other background characteristics that are correlated with caste. Consistent with survey responses of HR managers hiring for top MNCs in the U.S. as well as India, such background characteristics (e.g., family background, neighborhood of residence, and cultural fit) can plausibly only be learned during personal interviews and therefore cannot be used to differentiate candidates earlier (Section 6). However, this alternative interpretation of employers not caring about caste directly would be a bit at odds with the caste penalty emerging immediately during the application reading round, but only amongst those with obvious caste names (Figures 2 and 3). Still, even when employers "only" value caste-correlated traits, but not caste directly, such screening practices could create barriers to economic mobility for disadvantaged groups, especially as impressions of cultural fit are often skewed by socioeconomic considerations. If the current equilibrium in India is such that students from disadvantaged castes do not have the opportunity to develop the "right" soft skills or are perceived to not hail from the "right" socioeconomic background, they might face significant barriers in accessing top jobs.

Although my preferred interpretation of caste disparities in hiring is that they stem from employers directly penalizing caste, the magnitude of the caste penalty is *invariant* to the true interpretation or source of the caste penalty. Under a scenario where employers do not value caste at all, the "caste" coefficient in my regressions could be interpreted as a reduced-form capturing discrimination due to employers valuing other caste-correlated background characteristics learned by them during personal interviews. Caste disparities could stem from either taste-based or statistical discrimination and the reduced-form caste coefficient *would embed a mechanism for both.*

8.4.1 Back to Becker

My results beg the question of why the initial earnings gap across castes shown in Figure 1 is not competed away à la Becker (1957). Many social mechanisms could preclude unbiased firms from hiring disadvantaged castes or having disadvantaged castes form their own firms. Moreover, as documented at length in the California v. Cisco Systems, Inc. (2022) lawsuit, coworker discrimination against disadvantaged castes at elite multinationals is not merely a theoretical possibility. This paper does not attempt to conclusively label the type of discrimination operative in this setting, although my results do suggest that some labels (e.g., accurate statistical discrimination) might be less likely than others. The main goal of the paper is to *quantify* how much of the hiring disparities at elite jobs, however motivated, could be attributed to subjective assessments of "fit" done through personal interviews. Moreover, as stated in Section 8.4 above, the magnitude of the caste penalty is invariant to the true interpretation or source of the caste penalty.

8.4.2 Personal interviews are used to merely express a "delayed preference" for caste and are not necessarily a mechanism for caste revelation.

One alternative explanation for my preferred interpretation of the caste penalty could be that personal interviews are used to merely express a "delayed preference" for caste and are not necessarily a mechanism for caste revelation. I offer some arguments below for why this explanation is not of first-order importance in my setting.

First, even if caste disparities are expressed "later" or not at all (in the sense that employers only value other caste-correlated, background characteristics), this would not necessarily change the *magnitude* of the caste disparities in initial hiring. Indeed, as discussed in Section 8.4, none of my results crucially hinge on caste being revealed only at the personal interview round or the caste penalty in initial hiring stemming from employers caring about caste per se. Second, it remains unclear why employers do not penalize caste earlier if they already know it before personal interviews for a vast majority of the sample. As shown in Section 8.2, a caste penalty in the success rate emerges immediately at the application reading round, but only in a small sample of students with identifiable caste names. One might argue that, due to their subjective nature, the final screening round comprising informal personal interviews gives employers enough leeway, thereby allowing them to delay their preference for caste. However, it is unclear why employers cannot use group debates in a similar manner to engage in such dilatory methods. At group debates, employers presumably have as much free rein to assess social skills as they do when assessing fit during personal interviews. At least after the written test round, employers could have their hands tied because tests have objective scores attached to them. However, no such constraint is present during group debates that are conducted before personal interviews and are very subjective in their assessment of the "right" soft skills. Moreover, as will be discussed in Section 9.1 below, employers are not bumping up weak performers at earlier screening rounds to "appear fair" (ensure caste balance at each screening round). Neither is the caste penalty due to inframarginal differences in performance at the screening rounds prior to personal interviews that could be used by employers to remove the worst *cumulative* performers among those who qualify for personal interviews (Section 9.2).

9 Other Key Alternative Explanations

In Sections 9.1 and 9.2 below, I show that the caste penalty in hiring outcomes cannot be rationalized by jobs bumping up weaker performers before personal interviews to "appear fair" or due to large inframarginal caste differences in performance at screening rounds before personal interviews.

9.1 The Caste Penalty After Personal Interviews is Not Due to Jobs Bumping Up Weaker Performers Earlier to Appear Fair

One potential alternative theory is that caste disparities might be concentrated at personal interviews because jobs potentially lower the standard and bump up weaker performers at earlier screening rounds to "appear fair" (ensure caste balance at each screening round), only to later use personal interviews to "catch up" to previous performance and not extend job offers to such students. Recall, I do not observe direct performance measures at any screening round. Instead, I only observe if students qualified past (i.e., survived) a given screening round.

9.1.1 Quantitative Evidence: Disadvantaged Castes Do Not Have a Higher Success Rate than Advantaged Castes at Screening Rounds Prior to Personal Interviews Across the Full Distribution of Job Selectivity

I can quantitatively rule out a scenario where employers are lowering standards before interviews to appear fair. To do so, I use large variation across jobs in the cumulative proportion of students cut by the end of each screening round (defined as "cumulative cuts" in Section 5.3.1). Using this variation, I test if there is a "caste advantage" (i.e., disadvantaged castes having a higher success rate compared to advantaged castes) before the final screening round comprising personal interviews, especially among highly selective jobs.

The intuition for this test is straightforward. Consider a job that cuts a high cumulative proportion of applicants before personal interviews. Such a job would probably have to lower standards considerably to achieve caste balance at each screening round, especially since disadvantaged castes have, on average, weaker academic credentials than advantaged castes. If there are enough such jobs in the sample, then conditional on characteristics such as GPA, one should expect disadvantaged castes to eventually have a *higher* success rate compared to advantaged castes at screening rounds before personal interviews, especially among the most selective jobs. Essentially, variation in job selectivity is used to truncate the distribution of latent candidate talent and examine if a caste advantage arises in the success rate when one goes further and further to the right of this distribution as candidate screening advances. I formalize the results of this test through Figure 4. Specifically, for each job $j \in \mathcal{J}$, I estimate the following linear probability model:

$$\Pr(\text{Success}_{ijt}^{\text{Job Search Stage}}) = \alpha_j + \beta_j \times \text{Disadv. Caste}_i + \text{Controls}_i + \gamma_t + \text{error}, \quad (5)$$

where $i \in \mathcal{I}$ indexes students and γ_t denotes year fixed effects. The controls are standard characteristics such as college GPA, college entrance exam score, previous labor market experience, major, and degree. In Figure 4, I show the caste gap in the success rate with controls against the cumulative proportion of students cut by the end of each job search stage. Each blue triangle denotes a job. The x-axis shows the cumulative cuts made by jobs by the end of each screening round. Note that as we move from the top panel to the bottom panel of Figure 4, the lower and upper bounds of the x-axis increase because jobs cumulatively cut larger and larger proportions of their initial candidate pool as screening advances. The y-axis shows the conditional (i.e., with controls) caste gap in the success rate. A caste gap in the success rate below zero means that disadvantaged castes have a lower success rate or a higher cut rate than advantaged castes ("caste disadvantage").

I find no evidence of a notable caste advantage in the success rate at any of the three screening rounds prior to personal interviews, essentially across the *entire distribution* of job selectivity. This result suggests that jobs are not lowering standards to bump up weaker performers before personal interviews to appear fair. Panel (a) shows a negligible caste advantage in the conditional success rate at the application reading round across the distribution of job selectivity. Advancing forward from the application reading round, both panels (b) and (c) also show negligible caste advantage in the conditional success rates at the written test round and the group debate round, respectively. These results also hold across jobs with varying levels of selectivity. However, as shown in panel (d) of Figure 4, a substantial caste penalty in the success rate emerges at the fourth, and final, screening round comprising subjective, non-technical personal interviews, essentially across the *entire distribution* of job selectivity. Panel (d) shows that the average caste penalty (thick black line) is about 30 percentage points. This penalty is about the same as the one reported in Column (4) of Table 6 through linear probability regressions with job fixed effects. Overall, these results are inconsistent with jobs bumping up weaker performers before personal interviews to appear fair. In other words, there appears to be no differential scrutiny by caste for the majority of the student sample before personal interviews. A similar argument also rules out caste differences in the propensity to be at the margins of the unobserved "cutoff thresholds" employed by jobs to advance candidates at screening rounds prior to personal interviews. Recall, a caste penalty in the success rate emerged immediately at the application reading round, but only amongst those with obvious caste names (Figures 2 and 3). If employers were trying to appear fair or virtue signal, one would think that students they would most want to give a bump to are those with obvious caste names.

9.1.2 Qualitative Evidence: Weak Government or Internal Institutional Pressure to Promote Caste Diversity at Elite MNCs

The above quantitative evidence is also consistent with qualitative evidence suggesting little pressure on recruiters at elite MNCs to bump up disadvantaged castes at screening rounds prior to personal interviews to "appear fair." This lack of impetus is not only due to the absence of a federal ombudsman to regulate private-sector hiring in India but also due to weak internal institutional pressure at elite MNCs to promote caste diversity.

India lacks a formal regulatory agency, such as the Equal Employment Opportunity Commission (EEOC) in the United States. Informal government-led discussions have also found the private sector lacking support, even for the voluntary adoption of basic codes of conduct for advancing policies to promote caste diversity in hiring (Jodhka, 2008). Additionally, in the absence of an EEOC-like ombudsman, government pressure has been weak.

Overtly expressed attitudes by private-sector employers also suggest a lack of support for policies to promote caste diversity in hiring, even absent formal government pressure. For example, survey responses of HR managers at elite private-sector firms have found unanimous opposition toward consciously promoting caste diversity in hiring. Some respondents have even suggested that being caste-conscious is synonymous with being anti-merit (Jodhka and Newman, 2007; Newman and Thorat, 2012). Such opinions have been stated not only in private, but also publicly, with some insisting that hiring processes in the private sector are already caste-blind (Jodhka, 2008). Perhaps unsurprisingly, as recently as 2018, only 3 of the top 100 firms listed on India's premier stock exchange collected caste data for internal HR purposes (BusinessLine, 2018). Moreover, diversity practices in foreign-based multinationals employing Indians are overwhelmingly influenced by historical considerations of the West, where caste is not a protected category, suggesting that internal institutional pressure to rectify caste disparities have either been sluggish or, worse, non-existent (Chakravartty and Subramanian, 2021).

Finally, if internal institutional pressure on employers to "appear fair" is particularly pressing, it seems reasonable for these constraints to also be present at the job offer stage where, paradoxically, almost all caste disparities in hiring arise in my setting. Final hires are more visible to the broader public than shortlists from initial screening rounds.

9.2 The Caste Penalty is Not Due to Inframarginal Differences in Performance at Screening Rounds Prior to Personal Interviews

The arguments laid out in Section 9.1 show that there are negligible caste differences in the propensity to be at the margins of the unobserved "cutoff thresholds" employed by jobs to advance candidates at screening rounds before personal interviews. This may be because, conditional on qualifying for a screening round, everyone is comfortably above the cutoff threshold. However, there may still be large inframarginal caste differences in performance before personal interviews that could be used by employers to remove the worst *cumulative* performers among those who qualify for personal interviews.

To rule out such a scenario, I again use large variation across jobs in the cumulative proportion of students cut by the end of each screening round (defined as "cumulative cuts" in Section 5.3.1). If there are large caste differences in performance before personal interviews, then disadvantaged castes should eventually have a *lower* success rate than advantaged castes at the screening rounds before personal interviews, especially among the most selective jobs. However, I show that there is a negligible "caste disadvantage" (i.e., disadvantaged castes having a lower success rate compared to advantaged castes) before personal interviews, essentially across the *entire distribution* of job selectivity. Only after personal interviews do disadvantaged castes have a notably lower success rate than advantaged castes, essentially across the *entire distribution* of job selectivity (Figure 4). These results are inconsistent with employers relying on cumulative performance differences at screening rounds prior to personal interviews to advance disadvantaged castes at lower rates following personal interviews. If cumulative performance differences before personal interviews were significant, one should have seen a caste disadvantage in the success rate at screening rounds prior to personal interviews, especially among the most selective jobs.

9.2.1 Similar patterns of caste differences in the success rate at each screening round arise with or without other controls

Figure 5 shows a version of Figure 4 but with the y-axis replaced by the "raw" caste gaps in the success rate at each screening round. The "raw" caste gap in the success rate is calculated without the presence of any other controls.

In Figure 5, each blue triangle denotes a job. The x-axis shows the cumulative cuts made by jobs by the end of each screening round. Note that as we move from the top panel to the bottom panel of Figure 5, the lower and upper bounds of the x-axis increase because jobs cumulatively cut larger and larger proportions of the candidate pool as screening advances. The y-axis shows the raw caste gap in the success rate at each screening round. A raw caste gap in the success rate below zero means that disadvantaged castes have a lower raw success rate or a higher raw cut rate than advantaged castes.

Panel (a) shows that there is a raw caste penalty in the success rate at the application reading round. The pattern in panel (a) reflects the fact that disadvantaged castes have a stronger left tail on many employer-relevant characteristics such as GPA, and most students from the bottom-most end of the ability distribution get cut immediately when applications are read. Interestingly, when we move from panel (a) to panels (b) and (c), we see a successively smaller caste gap in the success rate at the written aptitude test round and the group debate round. In fact, even the *raw* caste gap in the success rate at the group debate round is negligible across the entire distribution of job selectivity. However, similar to Figure 4, a substantial caste penalty in the success rate emerges at the personal interview round, essentially across the *entire distribution* of job selectivity. Interestingly, the caste penalty shown in panel (d) of Figure 5 is similar in magnitude to that shown in panel (d) of Figure 4.

Note also that the number of blue triangles (jobs) shown in Figure 4 is slightly less than that shown in Figure 5. This is due to some jobs being very selective, making it empirically challenging to estimate a job-specific caste gap in the success rate with controls, especially at the personal interview round (see panel (d) of Figure 4). Still, the results shown in Figure 4 clearly support a caste penalty in the success rate emerging primarily after personal interviews across the *entire distribution* of job selectivity.

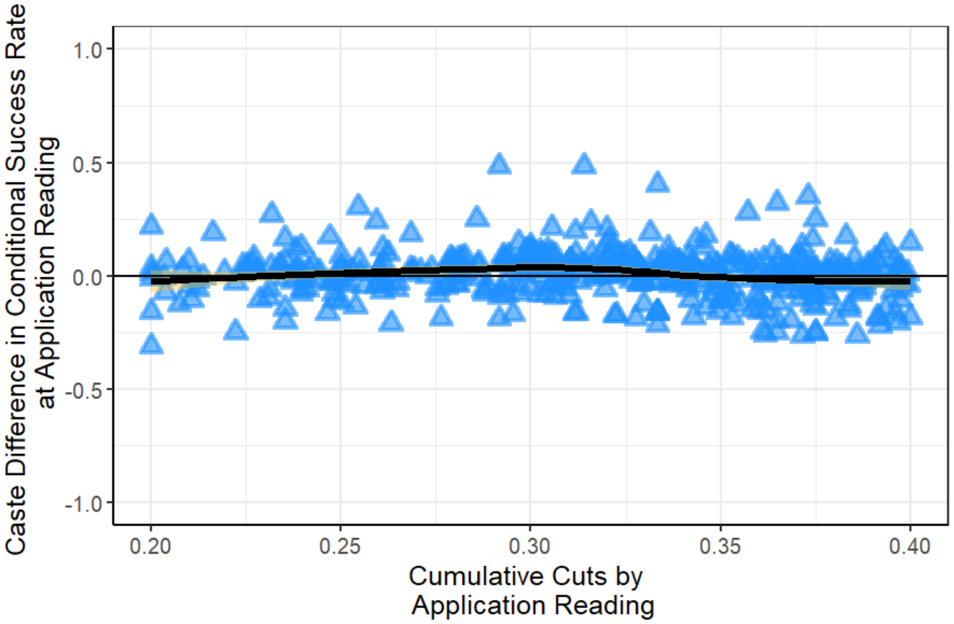

(a) Caste Gap in Conditional Success Rates versus Cumulative Cuts by Application Reading

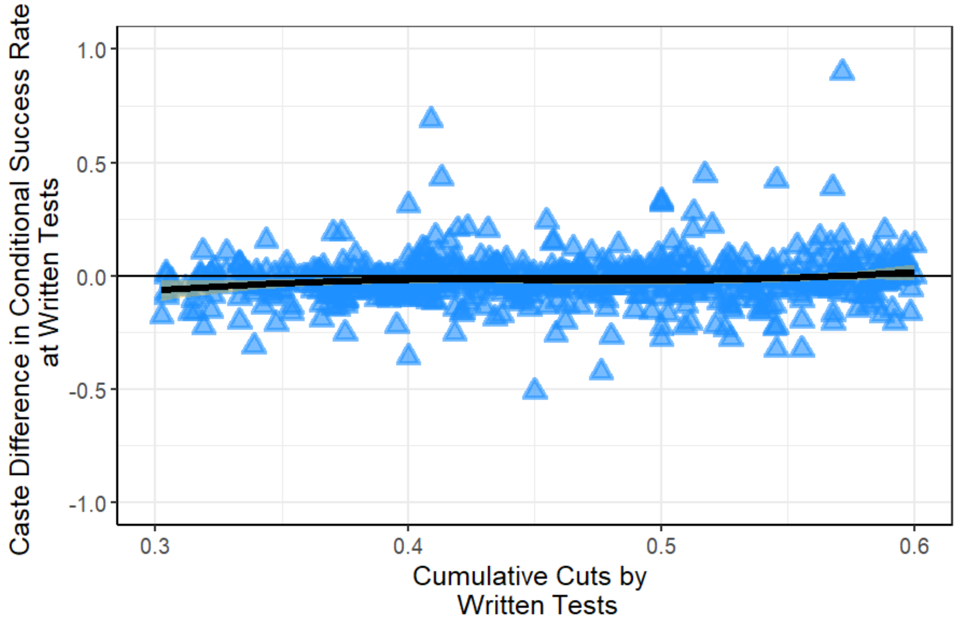

(b) Caste Gap in Conditional Success Rates versus Cumulative Cuts by Written Tests

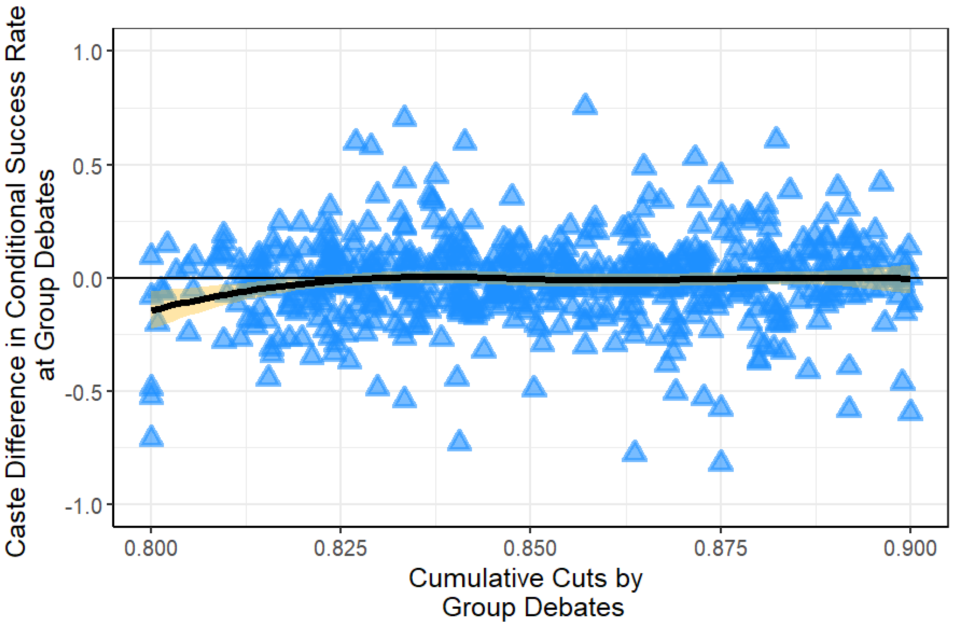

(c) Caste Gap in Conditional Success Rates versus Cumulative Cuts by Group Debates

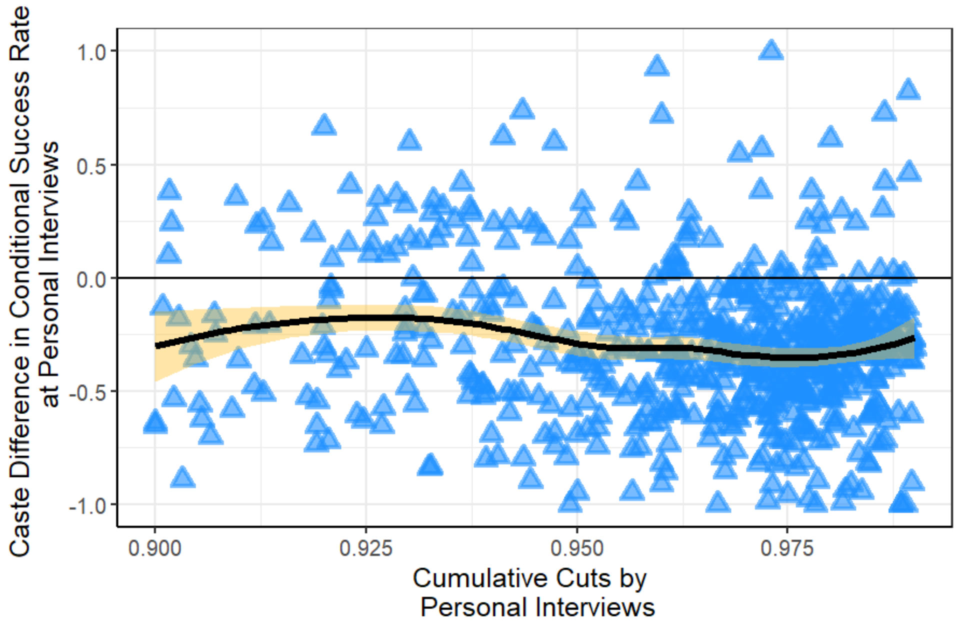

(d) Caste Gap in Conditional Success Rates versus Cumulative Cuts by Personal Interviews

Figure 4: Caste Penalty in the Conditional Success Rate At Each Screening Round

Notes: Figure 4 shows the caste difference in the conditional (i.e., with controls) success rate against the cumulative cuts made by jobs by the end of each screening round. Each blue triangle denotes a job. Smoothed conditional means (thick black lines) are calculated using the "loess" function in R. The shaded areas in yellow denote 95% confidence intervals. The thin solid horizontal line denotes a y-intercept of 0. The x-axis shows the cumulative proportion of students cut by jobs by the end of each screening round. The lower and upper bounds of the x-axis increases as one moves from panels (a) to (d) because jobs progressively cut more students from the initial candidate pool as screening advances. The y-values are calculated by estimating, for each job, a standard linear probability model with controls and then plotting the coefficient on disadvantaged caste from these regressions on the y-axis.

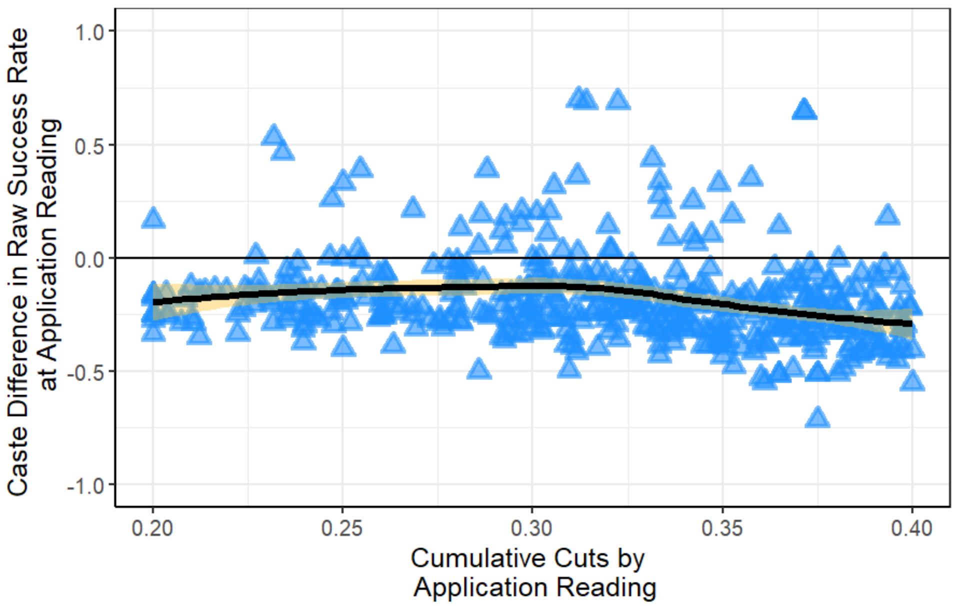

(a) Caste Gap in Raw Success Rates versus Cumulative Cuts by Application Reading

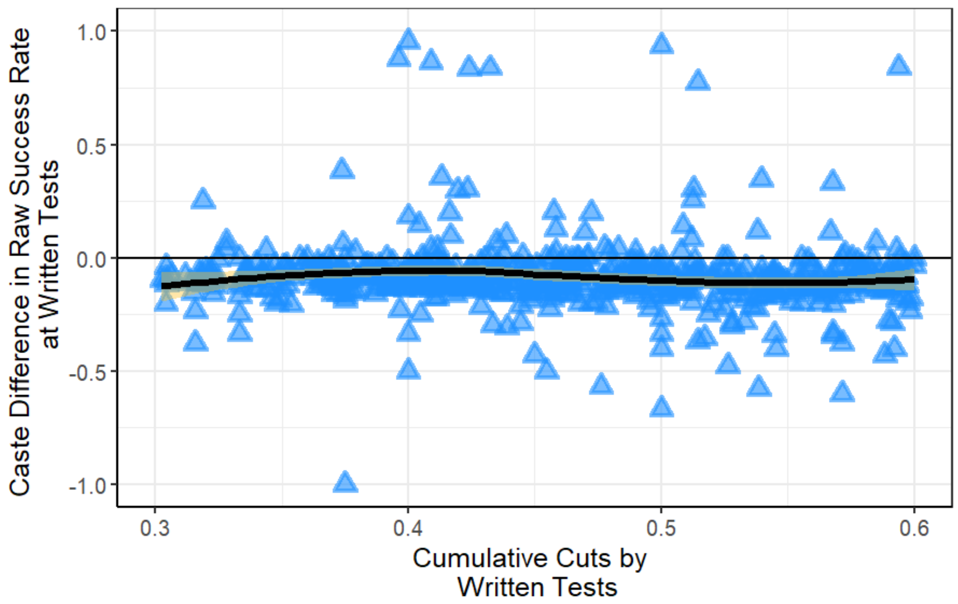

(b) Caste Gap in Raw Success Rates versus Cumulative Cuts by Written Tests

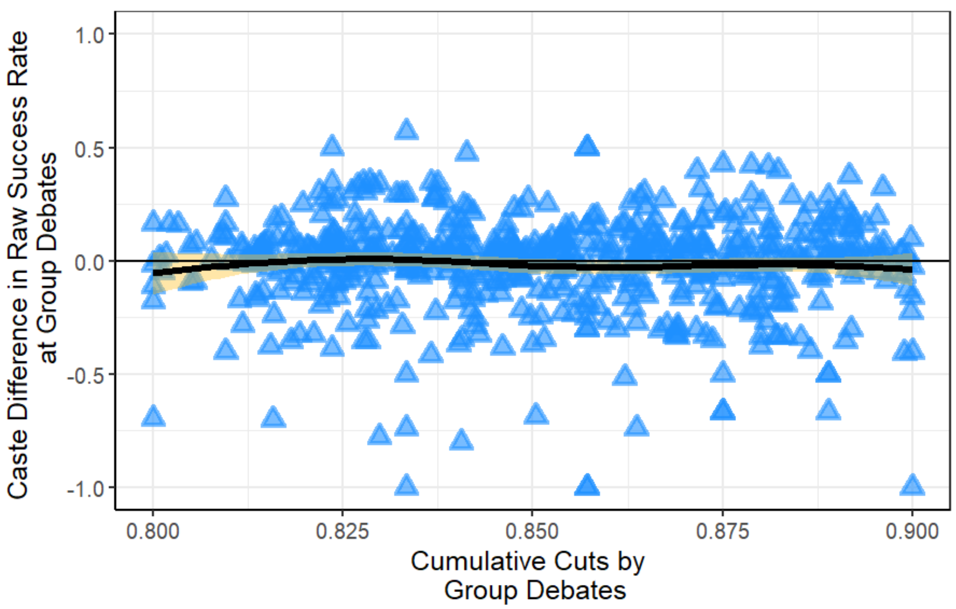

(c) Caste Gap in Raw Success Rates versus Cumulative Cuts by Group Debates

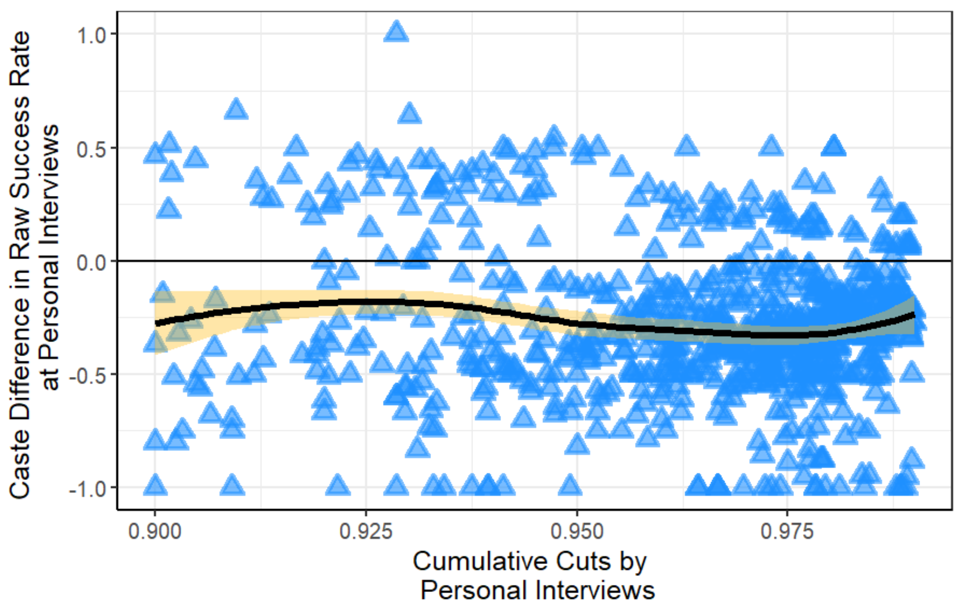

(d) Caste Gap in Raw Success Rates versus Cumulative Cuts by Personal Interviews

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

Notes: Figure 5 shows the caste difference in the "raw" (i.e., without any controls) success rate against the cumulative cuts made by jobs by the end of each screening round. Each blue triangle denotes a job. Smoothed conditional means (thick black lines) are calculated using the "loess" function in R. The shaded areas in yellow denote 95% confidence intervals. The thin solid horizontal line denotes a y-intercept of 0. The x-axis shows the cumulative proportion of students cut by jobs by the end of each screening round. The lower and upper bounds of the x-axis increases as one moves from panels (a) to (d) because jobs progressively cut more students from the initial candidate pool as screening advances. The y-axis shows the raw caste gap in the success rate at each screening round.

10 More Alternative Explanations

In this section, I discuss more alternative explanations that could explain caste disparities in hiring outcomes arising primarily after personal interviews.

1. **Lack of meaningful student attrition at each job search stage.** There are meaningful cuts being made by jobs at each screening round (Figures 4, 5, and Table 4).

2. **Jobs making offers before conducting all four screening rounds.** Every job in my sample conducts four screening rounds (application reading, written tests, group debates, and non-technical personal interviews) before making job offers (see Section 5.3.2). Likewise, having applied, students cannot "reject" firms midway or accept job offers "early" (see Section 4.4).

3. **Differences in socio-emotional skills.** This explanation is not consistent with panel (c) of Figures 4 and 5 that show negligible caste differences in the success rate at the group debate round among jobs with varying levels of selectivity. We would have seen a notable caste gap in the success rate below zero in panel (c) of Figures 4 and 5, if caste differences in socio-emotional skills were substantial.

4. **Better outside options for advantaged castes procured "offline."** In this elite college, students must decide on summer internship return offers *before* registering for the full-time job placement process. As already discussed, less than 5% of the graduating cohort deregisters from the placement process of full-time jobs by accepting summer internships (Section 4.3).

   Moreover, registered students are prohibited by the college from searching "offline" (i.e., outside of the centralized placement process) and, if discovered doing so, risk being debarred from the services of the placement office in assisting them during their full-time job search. Thus, registered students do not simultaneously search offline. As also mentioned in Section 4.3, it is beyond the scope of the paper to understand why certain placement rules exist. Such rules can be easily found on the websites of most elite Indian colleges.[36]

5. **Advantaged castes can bargain better over salaries and amenities.** Salaries and non-pecuniary amenities posted by jobs are non-negotiable during the course of the on-campus placement process (see "Pre-Placement phase" in Section 4.3). Compensation bundles are also verified by the placement office at the conclusion of the placement process. Furthermore, "exit data" confirms that job getters in my setting start in the same jobs and receive the same compensation bundles, months after the placement process has concluded (Section 7.1.1).

6. **Competition from the government sector.** While government jobs have caste-based quotas, they comprise less than 2% of all jobs in my sample (Section 5.2). Moreover, private-sector firms are themselves aware of the meager presence of public-sector firms in the on-campus job fair of the elite college. This is because the list of firms recruiting from campus is available to everyone on the online job portal designed by the college for the job placement process.

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

7. **Caste differences in preferences over job characteristics.** As suggested by Figure 1, differential caste preferences over job characteristics explain only a small portion of the overall caste penalty in initial hiring.

11 Policy Implications, in India and Elsewhere

The paper's descriptive evidence is helpful in identifying which policies are likely to be effective in remedying hiring disparities, in India and elsewhere. Potential policies need to be consistent with India's legal environment, which enforces explicit caste-based affirmative action policies, but only in public jobs and colleges. Therefore, unlike in the U.S., there is no federal ombudsman to regulate private-sector hiring. Moreover, while explicit caste-based discrimination is illegal, the Indian legal system does not recognize "disparate impact," and there is no systematic legal provision (anywhere) to penalize employers for judging "cultural fit" based on myriad characteristics correlated with protected status (Rivera, 2015; Jodkha, 2017; Lang and Spitzer, 2020). By quantifying the sources of caste disparities in elite jobs, policymakers could better understand the appropriate response in this context. Given that caste disparities primarily emerge after personal interviews in the multi-stage job search process, policies informing applicants about job opportunities, modifying student preferences, or improving performance at university are unlikely to meaningfully mitigate caste disparities. Therefore, the empirical quantification of the sources of caste disparities in elite hiring outcomes offers insights into better understanding the most effective policies to promote caste diversity among elite jobs hiring in India, nearly all of which hire entry-level employees from elite colleges such as the one in my setting (Section 4.2). Moreover, given potential productivity gains from hiring disadvantaged castes, elite multinationals hiring in India might even voluntarily adopt restrictions on the types of questions asked during informal interviews. Such policies also have relevance in the United States where, despite stronger employment laws, qualitative surveys have shown that the types of questions asked during open-ended, informal interviews tend to heavily favor the socioeconomically advantaged (Rivera, 2012, 2015). And even though the basis for hiring disparities in the United States is often "visual," reducing reliance on informal interviews could reduce socioeconomic and therefore racial or ethnic disparities in access to elite jobs. It is also not clear if hiring managers conducting informal evaluations have superior information or know what they are doing (Hoffman et al., 2018). While truly "blind" hiring à la Goldin and Rouse (2000) may not be possible in the context of elite job recruitment, reducing reliance on informal screening tools prone to bias is worth considering, in India and elsewhere.

Other policy avenues might also be promising. Policies with different types of "information provision" are typically proposed to reduce academic or career gaps across race or gender in the U.S. (Zafar and Wiswall, 2015a,b). In this setting, employers could be informed about the higher tenure and promotion rates of disadvantaged castes. The potential success of such an informational intervention would depend upon whether the discrimination is taste-based or statistical based on biased beliefs. Regardless, such a policy intervention may help further clarify the underlying mechanisms for the caste penalty in hiring.

Although no formal caste-based affirmative action policy exists in the Indian private sector, the government has also considered potential policy proposals to promote caste diversity. As mentioned in Section 9.1.2, the private sector has strongly opposed caste-quotas. Therefore, the government has made little to no progress on that front (Jodhka, 2008). As an alternative to hiring quotas, an incentive-based Diversity Index to promote caste diversity in the private sector was proposed by the Indian Ministry of Minority

Affairs in 2008 (Sachar Committee, 2006; Report of the Expert Group on Diversity Index, 2008):

> The implementation of the [diversity] index prepared by this Expert Group could be operationalised either on the principle of (a) incentive (reward) or disincentive (punishment) or (b) a lucrative incentive(s) (reward) and the lack of it (the denial of reward to be construed as penalty). Obviously, the former could be resented, may even lead to legal entanglements, while the latter, though slower to implement, could initially be used by enthusiastic States and institutions for incentives, while the others may just ignore it. But eventually, the Expert Group feels, it would catch up.

A subsidy to hire disadvantaged castes could be a natural way to operationalize this proposal. Yet, it is possible that subsidies could stigmatize its beneficiaries (Coate and Loury, 1993). However, empirical research both in India and the U.S. has found slim evidence in support of *further* stigmatization of minorities due to affirmative action policies (Bowen and Bok, 1998; Deshpande, 2018). Further, even granting the purported worsening of stigma, such policies could still be efficiency enhancing, as disadvantaged groups likely benefit the most from elite attainment, whereas displaced advantaged groups are likely not much worse off (Black et al., 2020). The last point is also why it could still be meaningful to intervene through policies to address disparities from "accurate" statistical discrimination.

12 Limitations of the Study

In addition to the data limitations discussed in Section 5.6, one of the main limitations of this study is that it focuses on hiring from only one elite Indian college and only among elite-entry level jobs in the Indian private sector. I argue below why these potential limitations do not necessarily restrict the scope of this study. Jobs in the Indian private sector pay much higher than those in the public sector and thus are plausibly more relevant to the economic mobility of historically marginalized groups. Therefore, such jobs are important to focus on as they can shape not just an individual's economic trajectory but also broader societal inequalities (Rivera, 2012; Subramanian, 2019). Moreover, while the jobs in my setting hire for their Indian locations, they are located at elite multinationals primarily headquartered in the U.S. or Europe. Thus, the recruitment practices I study are likely not limited only to India. Indeed, personal interviews that assess fit are a common component of elite job recruitment at American workplaces (Rivera, 2012, 2015). Additionally, the job placement process of this elite college offers a representative window into how elite college graduates transition into elite entry-level jobs in the Indian private sector. This is essentially because nearly all elite Indian colleges conduct similar job fairs with participation from nearly the same group of firms (see Section 4.2). Moreover, several other studies in fields ranging from sociology, psychology, organizational decision making, and medicine have shown that unstructured interviews, evaluations, or other forms of subjective assessments are common avenues where disparities emerge. Thus, my core finding that personal interviews could lead to unwarranted disparities has been replicated across multiple fields. However, this paper goes beyond showing that personal interviews are common avenues where disparities emerge, quantifying the real effects of disparities due to such subjective screening practices in earnings and access to elite jobs. To my knowledge, this paper is the first to quantify how much of the hiring disparities among elite jobs in a real world and high-stakes setting, could be attributed to subjective assessments of fit done through personal interviews (see Section 7.6.3).

13 Conclusion

"What are your favorite hobbies?," "Where do you live?," and "Where are you from?" are real questions asked during candidate screening by HR managers at elite MNCs hiring in India, alumni interviewers at the Ivy League universities, and hiring professionals at top U.S. firms (Jodhka and Newman, 2007; Stevens, 2009; Rivera, 2012, 2015).

How quantitatively important are screening practices in driving socioeconomic gaps in access to elite jobs, and what aspects of screening generate such disparities? The lagging representation of disadvantaged castes in elite jobs, despite broad-based affirmative action in university admissions in India, could be partly driven by screening processes which allow employers to discriminate against disadvantaged castes. However, little is known about how *quantitatively* important such screening processes are in driving caste disparities in access to the "elite," or *what* aspects of candidate screening drive these disparities. Adding to these challenges is the fact that recruitment practices of elite jobs are often non-transparent (Rivera, 2015). My paper addresses these issues by showing that the final screening round consisting of subjective, non-technical personal interviews that often assess "fit" could create barriers for disadvantaged castes in accessing elite jobs. Specifically, I show that nearly 90% of the caste disparities in hiring arise due to personal interviews. I also argue that this screening round is where employers are most likely to learn about caste. Further, information learned during personal interviews is weakly correlated with job performance. Instead, personal interviews, a common component of elite job recruitment in India and elsewhere, are a large contributor to unwarranted socioeconomic (caste) disparities in the labor market. Such disparities could be motivated by taste-based, statistical, customer, or coworker discrimination, or other factors. To my knowledge, this paper is the first to quantify how much of the hiring disparities, however motivated, among elite jobs in a real world and high-stakes setting, could be attributed to subjective assessments of "fit" done through personal interviews.

While this paper attempts to do many things, no paper is exhaustive. As with any study that uses hiring data from one elite college, there should be further research to establish more external validity. Studying elite recruitment in different settings might provide novel insights or yield results different from those presented in this paper. Future research could also collect similar data to detect less visible or covert forms of discrimination in other parts of the world. Experimental follow-ups studying different firm-level policies such as standardized interview questions, decision review, and interviewer representation are also promising areas for future exploration. Examining the evolution of the caste penalty beyond the first job could also be promising.

Disparities based on socioeconomic cues are likely to become more salient in elite, urban-educated settings because standard characteristics by which to differentiate groups become less perceptible as these settings become more multi-ethnic and diverse (Loury, 2002; Freeman et al., 2011; Gaddis, 2017). I provide an important example of the kinds of data future researchers may have to collect in such settings to better detect disparities from outwardly neutral screening practices. Even in much of the Western world, as racial boundaries blur, "caste" discrimination is likely to become a bigger part of the discourse globally (Austen-Smith and Fryer, 2005; Eguia, 2017; Kim and Loury, 2019; Rose, 2022). By connecting how perceptions of socioeconomic cues determine barriers to economic mobility and elite attainment, this paper also helps advance how to conceptualize, quantify, and address racial, class, or caste disparities in such opportunities, most of which are situated in a rapidly diversifying urban landscape.

References

[1] **Altbach, Philip G.**, *A Half-Century of Indian Higher Education*, SAGE Publications, 2012.

[2] **Anthropological Survey of India**, "People of India (Bihar & Jharkhand Project)," 2009. Ministry of Culture, Government of India.

[3] **Anzia, Sarah F. and Christopher R. Berry**, "The Jackie (and Jill) Robinson Effect: Why Do Congresswomen Outperform Congressmen?," *American Journal of Political Science*, 2011, *55* (3), 478–493.

[4] **Arcidiacono, Peter, Josh Kinsler, and Tyler Ransom**, "Legacy and Athlete Preferences at Harvard," 2021. Forthcoming, *Journal of Labor Economics*.

[5] **Austen-Smith, D. and Ronland Fryer**, "An Economic Analysis of "Acting White"," *Quarterly Journal of Economics*, 2005, *120* (2), 551–58.

[6] **Ayers, Ian**, "Outcome Tests of Racial Disparities in Police Practices," *Justice Research and Policy*, 2002, *4* (1-2), 131–142.

[7] **Ayres, Ian**, "Testing for Discrimination and the Problem of "Included Variable Bias"," 2010. Working Paper.

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

[9] **Becker, G.**, *The Economics of Discrimination*, Chicago: The University of Chicago Press, 1957.

[10] **Benson, Alan, Danielle Li, and Kelly Shue**, ""Potential" and the Gender Promotion Gap," 2023. R&R, *American Economic Review*.

[11] **Bertrand, Marianne, Abhijit Banerjee, Saugato Datta, and Sendhil Mullainathan**, "Labor Market Discrimination in Delhi: Evidence From a Field Experiment," *Journal of Comparative Economics*, 2009, *31* (1), 14–27.

[12] ___ **and Esther Duflo**, "Field Experiments on Discrimination," 2016. NBER Working Paper 22014.

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

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

[15] **Bhavnani, Rikhil R. and Alexander Lee**, "Does Affirmative Action Worsen Bureaucratic Performance? Evidence from the Indian Administrative Service," *American Journal of Political Science*, 2021, *65(1)* (5), 5–20.

[16] **Black, Sandra E., Jeffrey T. Denning, and J. Rothstein**, "Winners and Losers? The Effect of Gaining and Losing Access to Selective Colleges on Education and Labor Market Outcomes," 2020. R&R, *American Economic Journal: Applied Economics*.

[17] **Bowen, W. G. and Derek Bok**, *The Shape of the River*, Princeton University Press, 1998.

[18] **BusinessLine**, "An Inclusive India Inc.," 2018. Business Line, January 13, 2018.

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

[20] **Carroll, J. S., R. L. Wiener, D. Coates, J. Galegher, and J. J. Alibrio**, "Evaluation, Diagnosis, and Prediction in Parole Decision Making," *Law & Society Review*, 1982, *17*, 199–228.

[21] **Chakravartty, P. and Ajantha Subramanian**, "Why is Caste Inequality Still Legal in America?,"

2021. *The New York Times*, May 25, 2021.

[22] **Chetty, Raj, David J. Deming, and John N. Friedman**, "Diversifying Society's Leaders? The Determinants and Causal Effects of Admission to Highly Selective Private Colleges," 2023. NBER Working Paper 31492.

[23] **Coate, S. and Glenn C. Loury**, "Will Affrmative Action Eliminate Negative Stereotypes?," *American Economic Review*, 1993, *83*, 1220–1240.

[24] **Dana, Jason, Robyn Dawes, and Nathanial Peterson**, "Belief in the Unstructured Interview: The Persistence of an Illusion," *Judgement and Decision Making*, 2013, *8* (5), 512–520.

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

[26] **Dayanandan, Ajit., Han Donker, and Ravi Saxena**, "The Role of Caste for Board Membership, CEO, and Interlocking," *Pacific-Basin Finance Journal*, 2019, *54*, 29–41.

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

[28] ——, *Stigma or Red Tape? Roadblocks in the Use of Affirmative Action*, *The Empire of Disgust: Prejudice, Discrimination, and Policy in India and the US*, ed. by Zoya Hasan, 2018.

[29] —— **and Katherine Newman**, "Where the Path Leads: The Role of Caste in Post-University Employment Expectations," *Economic and Political Weekly*, 2007, *42* (41), 4133–4140.

[30] **DeVaul, R., F. Jervey, J. Chappell, P. Caver, B. Short, and S. O'Keefe**, "Medical School Performance of Initially Rejected Students," *Journal of the American Medical Association*, 1987, *257*, 47–51.

[31] **Eguia, Jon X.**, "Discrimination and Assimilation At School," *Journal of Public Economics*, 2017, *156*, 48–58.

[32] **Exit Survey of MBA Graduates, Yale School of Management**, "Yale SOM Employment Data," 2021. Accessed July 26, 2022.

[33] **Fernandez, Marilyn**, *The New Frontier: Merit vs. Caste in the Indian IT Sector*, Oxford: Oxford University Press, 2018.

[34] **Ferreira, Fernando and Joseph Gyourko**, "Does Gender Matter for Political Leadership? The Case of U.S. Mayors," *Journal of Public Economics*, 2014, *112*, 24–39.

[35] **Freeman, Jonathan B., Andrew M. Penner, Aliya Saperstein, Matthias Scheutz, and Nalini Ambady**, "Looking the Part: Social Status Cues Shape Race Perception," *PLOS ONE*, 2011, *6* (9), 1–10.

[36] **Frisancho Robles, Veronica C. and Kala Krishna**, "Affirmative Action in Higher Education in India: Targeting, Catch Up, and Mismatch," *Higher Education*, 2015.

[37] **Gaddis, S. Michael**, "How Black Are Lakisha and Jamal? Racial Perceptions from Names Used in Correspondence Audit Studies," *Sociological Science*, 2017, *4* (19), 469–489.

[38] **Goldin, Claudia and Cecilia Rouse**, "Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians," *American Economic Review*, 2000, *90* (4), 715–741.

[39] **Gumperz, John J.**, "Speech Variation and the Study of Indian Civilization," *American Anthropologist*, 1961, *63* (5), 973–988.

[40] **Hazell, Jonathan, Christina Patterson, Heather Sarsons, and Bledi Tsaka**, "National Wage Setting," 2021. Working Paper.

[41] **Hengel, Erin and Eunyoung Moon**, "Gender and Equality At Top Economics Journals," 2023. Work-

ing Paper.

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

[43] **Hoffman, Mitchell, Lisa B. Kahn, and Danielle Li**, "Discretion in Hiring," *The Quarterly Journal of Economics*, 2018, *133* (2), 765–800.

[44] **Hoxby, Caroline and Christopher Avery**, "The Missing 'On-Offs': The Hidden Supply of High-Achieving, Low Income Students," *Brookings Papers on Economic Activity*, 2013, (1), 1–65.

[45] **Hull, Peter**, "What Marginal Outcome Tests Can Tell Us About Racially Biased Decision-Making," 2021. NBER Working Paper 28503.

[46] **____ , Aislinn J. Bohren, and Alex Imas**, "Systemic Discrimination: Theory and Measurement," 2022. NBER Working Paper 29820.

[47] **Jayaraman, R.**, "Personal Identity in a Globalized World: Cultural Roots of Hindu Personal Names and Surnames," *The Journal of Popular Culture*, 2005, *38* (3), 476–490.

[48] **Jodhka, Surinder S**, "Caste & Corporate Sector," *Indian Journal of Industrial Relations*, 2008, *44* (2), 185–193.

[49] **Jodhka, Surinder S. and Jules Naudet**, *Mapping the Elite: Power, Privilege, and Inequality*, Oxford University Press, 2019.

[50] **____ and Katherine Newman**, "In the Name of Globalisation: Meritocracy, Productivity and the Hidden Language of Caste," *Economic and Political Weekly*, 2007, *42* (41), 4125–4132.

[51] **Jodkha, Surinder S.**, *Caste in Contemporary India*, Routledge India, 2017.

[52] **Kim, Young-Chul and Glenn C. Loury**, "To Be, Or Not To Be: Stereotypes, Identity choice and Group Inequality," *Journal of Public Economics*, 2019, *174*, 36–52.

[53] **Kline, Patrick, Evan Rose, and Christopher Walters**, "Systemic Discrimination Among Large U.S. Employers," *Quarterly Journal of Economics*, 2022, *137* (4), 1963–2036.

[54] **Kothari, Rima**, "Caste in a Casteless Language? English as a Language of 'Dalit' Expression," *Economic and Political Weekly*, 2013, *48* (39), 60–68.

[55] **Lang, Kevin and Ariella Kahn-Lang Spitzer**, "How Discrimination and Bias Shape Outcomes," *Future of Children*, 2020, *30* (1), 165–190.

[56] **Loury, Glenn C.**, "The Anatomy of Racial Inequality," 2002. Harvard University Press.

[57] **Mamgain, Rajendra P.**, *Formal Labour Market in Urban India: Job Search, Hiring Practices and Discrimination*, SAGE Publications Pvt. Ltd, 2019.

[58] **Mamidi, Pavan**, "Strategic Manipulation of Caste Identity Signals," 2011. Doctoral Dissertation, Oxford University.

[59] **Mandal Commission Report**, "Mandal Commission Report," 1980. New Delhi: National Commission for Backward Classes, Government of India. November 1, 1980. Retrieved March 26, 2019.

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

[61] **Mishra, Neha**, "India and Colorism: The Finer Nuances," *Washington University Global Studies Law Review*, 2015, *14*, 725–750.

[62] **Mogstad, Magne, Ivan Canay, and Jack Mountjoy**, "On the Use of Outcome Tests for Detecting Bias in Decision Making," 2023. Forthcoming in the *Review of Economic Studies*.

[63] **Naukri.com**, "10 Common Interview Questions & Answers for Freshers," 2019. Naukri.com, October

9, 2019.

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

[65] **Parameswaran, R. and K. Cardoza**, "Melanin on the Margins: Advertising and the Cultural Politics of Fair/Light/White Beauty in India," *Journalism & Communication Monographs*, 2015, *11* (3), 213–274.

[66] **Ransubhe, Suryanarayan**, *Sociology of Translation*, Atulya Publications, 2018.

[67] **Report of the Expert Group on Diversity Index**, "Report of the Expert Group to Propose a Diversity Index and Workout the Modalities of Implementation," 2008. Submitted to Ministry of Minority Affairs Government of India.

[68] **Rivera, Lauren A.**, "Hiring as Cultural Matching: The Case of Elite Professional Service Firms," *American Sociological Review*, 2012, *77* (6), 999–1022.

[69] ——, *Pedigree: How Elite Students Get Elite Jobs*, Princeton University Press, 2015.

[70] **Rose, Evan K.**, "A Constructivist's Perspective on Empirical Discrimination Research," 2022. Forthcoming, *Journal of Economic Literature*.

[71] **Sachar Committee**, "Social, Economic and Educational Status of the Muslim Community," 2006. Prime Minister's High Level Committee, Cabinet Secretariat, Government of India.

[72] **Sarsons, Heather, Jonas Hjort, and Xuan Li**, "Across-Country Wage Compression in Multinationals," 2022. R&R, *American Economic Review*.

[73] **Savage, Mike**, *Social Class in the 21st Century*, Pelican Books, 2015.

[74] **SFFA v. Harvard**, "SFFA v. Harvard," 2023. 600 U.S. 181.

[75] **Stevens, Michael P.**, *Creating a Class: College Admissions and the Education of Elites*, Harvard University Press, 2009.

[76] **Subramanian, Ajantha**, *The Caste Of Merit: Enginnering Education in India*, Harvard University Press, 2019.

[77] **The State of Inequality in India Report**, "The State of Inequality in India Report," 2022. Economic Advisory Council to the Prime Minister, Government of India.

[78] **Wiltshire, Caroline R.**, *Uniformity and Variability in the Indian English Accent*, Cambridge University Press, 2020.

[79] **Zafar, Basit and Matthew Wiswall**, "Determinants of College Major Choice: Identification using an Information Experiment," *The Review of Economic Studies*, 2015, *82* (2), 791–824.

[80] —— **and** ——, "How Do College Students Respond to Public Information about Earnings? ," *Journal of Human Capital*, 2015, *9* (2), 117–169.

[81] **Zimmerman, Seth**, "Elite Colleges and Upward Mobility to Top Jobs and Incomes," *American Economic Review*, 2019, *109* (1), 1–47.

A ONLINE APPENDIX

A.1 Representativeness of the On-Campus Job Fair of the Elite College

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

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

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

3. **Elite entry-level jobs almost exclusively hire from elite colleges.** See Subramanian (2019) and Jodhka and Naudet (2019) for a comprehensive overview.

4. **Small proportions of students skip the job fairs of elite colleges.** As mentioned in Section 4.3, students must register for the formal on-campus job fairs of elite colleges, but otherwise must "deregister" or formally drop out. Only a small proportion (less than 5%) of college graduates in my setting skip the formal placement process for full-time jobs (see Section 7.1.2). Given the similarity of the placement process across elite Indian colleges, this proportion is likely to be similar in other institutions as well. See Mamgain (2019) and Subramanian (2019) for more details.

   The short discussion in points 1), 2), 3), and 4) above imply that the job placement process from the college I examine in the paper offers a representative window into how elite college graduates transition into elite entry-level jobs in the Indian private sector.

A.2 Tables and Figures

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

	Dependent variable: log (Salary)			
	Without Controls	With Controls	Client Facing	Non-Client Facing
	(1)	(2)	(3)	(4)
Disadv. Caste	-0.231^{***}	-0.103^{***}	-0.107^{***}	-0.089^{***}
	(0.020)	(0.017)	(0.027)	(0.022)
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

Notes: Online Appendix Table OA.1 includes estimates from an earnings regression run on the sample of all students who graduated with jobs obtained through the college's on-campus job fair for full-time jobs. The dependent variable is log earnings. The regressions include detailed controls including measures of pre-college skills, college academic performance, previous labor market experience, and other employer-relevant skills. Column (2) reports estimates from a linear earnings regression without controls. Column (2) reports estimates from a linear earnings regression with controls. Columns (3) and (4) report estimates from a linear earnings regression with controls run separately for client-facing and non-client-facing jobs, respectively. $*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

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

	Dependent variable: log (Salary)					
	App	Test	Debate	Interview	Offer	Choice
	(1)	(2)	(3)	(4)	(5)	(6)
Disadv. Caste	0.004	-0.009	-0.006	-0.018^{*}	-0.105^{***}	-0.103^{***}
	(0.005)	(0.007)	(0.007)	(0.010)	(0.013)	(0.017)
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

Notes: Online Appendix Table OA.2 includes estimates from an earnings regression run on the sample of all students who remain in contention for some job at a given job search stage. For example, column (1) includes the sample of all students who apply for jobs recruiting from the college's job fair. Similarly, column (2) includes the sample of all students who remain in contention for some job after the application reading round (i.e., advanced to the next written test round) and so on. The dependent variable is log earnings. I include detailed controls including measures of pre-college skills, college academic performance, previous labor market experience, and other employer-relevant skills. Each column is a separate regression and includes these controls. $*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

Table OA.3: Comparing College GPA of Students Who Secured Full-Time Jobs Through the College's Job Fair Versus Not

	With Jobs		Without Jobs		Difference	
College GPA	Adv. Caste	Disadv. Caste	Adv. Caste	Disadv. Caste	Adv. Caste	Disadv. Caste
B.Tech.	0.87	0.50	-0.28	-0.89	-1.15***	-1.39***
Dual Degree	0.76	0.52	-0.41	-1.01	-1.17***	-1.53***
M.Tech.	0.60	0.35	-0.67	-1.17	-1.27***	-1.52***
M.S.	0.08	-0.04	-0.34	-0.64	-0.42**	-0.60**

Notes: Online Appendix Table OA.3 compares the average college GPA of students who graduated with full-time jobs obtained through the college's job fair versus not. College GPA is normalized to have zero mean and unit standard deviation within each college degree and year. The differences are reported in standard deviation units and are calculated within each caste. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

Table OA.4: Caste Differences in the Success Rate at Each Job Search Stage Among Client-Facing Jobs

	Success App (1)	Success Test (2)	Success Debate (3)	Success Interview (4)
Disadv. Caste	-0.0066***	-0.0007	-0.0136***	-0.3062***
	(0.0022)	(0.0030)	(0.0046)	(0.0081)
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

Notes: Online Appendix Table OA.4 shows the caste difference in the success rate at each stage of job search only among client-facing jobs. Each observation is a student who remains in contention at a given screening round in a given job in a given year. Each column reports a separate linear probability regression and includes controls. Note that some firms conduct the same screening round (written tests, group debates, interviews) for multiple job positions. However, since these job positions within the same firm represent different jobs, the regressions consider these jobs and the corresponding advance-or-out decisions made by them separately. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

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

	Success App (1)	Success Test (2)	Success Debate (3)	Success Interview (4)
Disadv. Caste	0.0024	0.0011	0.0070**	-0.2867***
	(0.0018)	(0.0022)	(0.0034)	(0.0060)
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

Notes: Online Appendix Table OA.5 shows the caste difference in the success rate at each stage of job search only among non-client-facing jobs. Each observation is a student who remains in contention at a given screening round in a given job in a given year. Each column reports a separate linear probability regression and includes controls that enter linearly. Note that some firms conduct the same screening round (written tests, group debates, interviews) for multiple job positions. However, since these job positions within the same firm represent different jobs, the regressions consider these jobs and the corresponding advance-or-out decisions made by them separately. $*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

Table OA.6: College GPA Comparison of Students Who Stayed versus Left First Job

	Stayed At First Job		Left First Job		Difference	
College GPA	Adv. Caste	Disadv. Caste	Adv. Caste	Disadv. Caste	Adv. Caste	Disadv. Caste
B.Tech.	1.04	0.54	0.67	0.45	-0.37***	-0.09*
Dual Degree	0.94	0.61	0.63	0.37	-0.31***	-0.24***
M.Tech.	0.78	0.46	0.41	0.19	-0.37***	-0.27***
M.S.	0.38	0.09	-0.20	-0.34	-0.58***	-0.43*

Notes: Online Appendix Table OA.6 compares the average college GPA of students who left their first job at some point during the timeframe for which promotion and tenure data is collected versus not. College GPA is normalized to have zero mean and unit standard deviation within each college degree and year. The differences are reported in standard deviation units and are calculated within each caste. $*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

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

	Dependent variable:					
	log (Salary)					
	App	Test	Debate	Interview	Offer	Choice
	(1)	(2)	(3)	(4)	(5)	(6)
Disadv. Caste	0.012	-0.068^{***}	-0.034^{*}	-0.105^{***}	-0.146^{***}	-0.107^{**}
	(0.014)	(0.019)	(0.018)	(0.027)	(0.033)	(0.049)
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

Notes: Online Appendix Table OA.7 includes estimates from an earnings regression run on the sample of students with identifiable caste names who remain in contention for some job at a given job search stage. For example, column (1) includes the sample of all students with identifiable caste names who apply for jobs recruiting from the college's job fair. Similarly, column (2) includes the sample of all students with identifiable caste names who remain in contention for some job at the written test round, and so on. The dependent variable is log earnings. Each column is a separate regression and includes controls. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

Table OA.8: Caste Differences in Mean Salary at Each Job Search Stage in the Nonidentifiable Sample

	Dependent variable:					
	log (Salary)					
	App	Test	Debate	Interview	Offer	Choice
	(1)	(2)	(3)	(4)	(5)	(6)
Disadv. Caste	0.002	0.0004	-0.003	-0.004	-0.100^{***}	-0.106^{***}
	(0.005)	(0.007)	(0.007)	(0.011)	(0.014)	(0.019)
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

Notes: Online Appendix Table OA.8 includes estimates from an earnings regression run on the sample of students with nonidentifiable caste names who remain in contention for some job at a given job search stage. For example, column (1) includes the sample of all students with nonidentifiable caste names who apply for jobs recruiting from the college's job fair. Similarly, column (2) includes the sample of all students with nonidentifiable caste names who remain in contention for some job at the written test round, and so on. The dependent variable is log earnings. Each column is a separate regression and includes controls. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

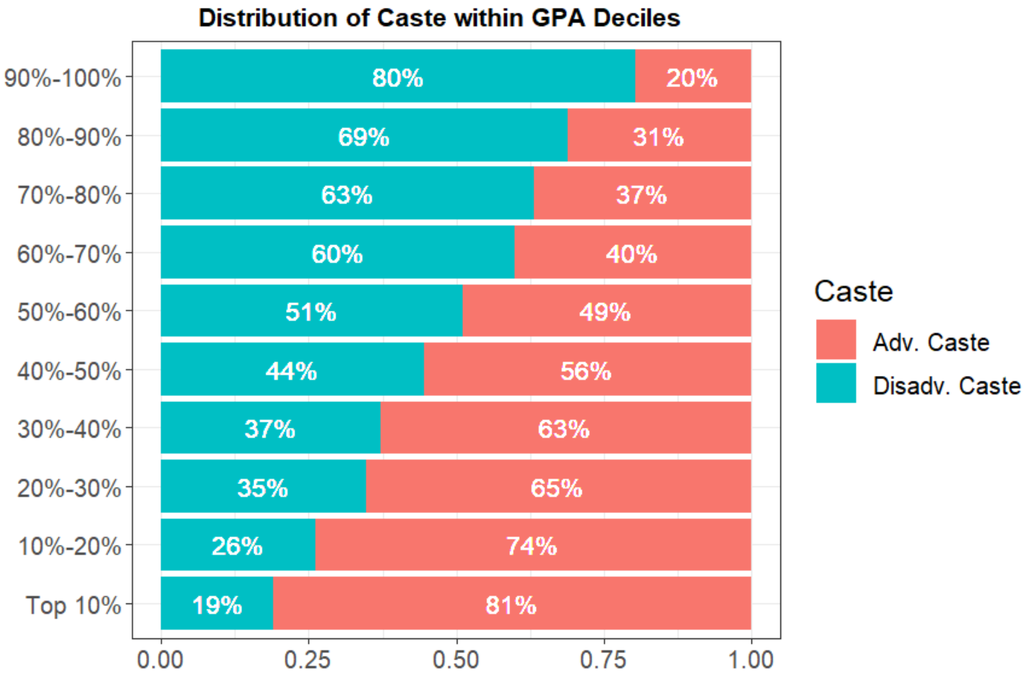

Figure OA.1: Common Support Within Each GPA Decile

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