Familiar strangers:
Evidence from referral-based hiring experiments in India

Kartik Srivastava*†

March 20, 2025

Job Market Paper

Preliminary draft for conference submissions. Please do not cite or circulate.

Click here for the latest version.

ABSTRACT. Across developing countries, search and matching frictions in labor markets are associated with high turnover rates in entry-level jobs. At the same time, large fractions of the workforce remain without access to certain jobs. Firms often rely on referral-based hiring in these settings to reduce search costs and improve match quality. However, these practices can systematically exclude underrepresented groups and exacerbate inequalities. This paper asks whether reallocating referral opportunities toward historically disadvantaged groups can improve firm performance and worker welfare, focusing on the role of caste in India. In a manufacturing firm in India, I experimentally increase the probability of referral opportunities being shared with lower caste incumbents. I find that this increases their representation by 15 percentage points (62%), and reduces turnover by 4 percentage points (41%) owing to worse outside options for these entrants. Output improves by 0.07 sd, and in contrast to the litearture on the adverse effects of mixed teams on cohesion, I find no significant effects on this in my setting. A lab-in-field experiment establishes a mechanism driving the results: cohesion is lower by 9% when lower caste workers are hired as outsiders without links to incumbent workers, which reduces total team output by 22% relative to when these workers are hired through referrals.

*Harvard University; kartiksrivastava@g.harvard.edu; Personal Website.

†I acknowledge with gratitude support from The Weiss Fund for Research in Development Economics and The Mittal South Asia Institute at Harvard University. I am especially grateful to Emily Breza, Asim Ijaz Khwaja, and Eliana La Ferrara for their invaluable feedback and guidance through every stage of the project. I am also grateful for helpful comments from Edward Glaeser, Rema Hanna, Gordon Hanson, Amen Jalal, Namrata Kala, Gabriel Kreindler, MR Sharan, and the participants of the Development, Labor, and Economic and Social Policy workshops at Harvard University. I thank Aiushee Bhati, Javed Khan and Sunil Kumar for excellent research assistance. This project is registered at the AEA Registry (AEARCTR-0013091), and received approval from the Harvard University IRB (IRB23-0323).

1 Introduction

An estimated 125 million people in low- and middle-income countries were unemployed in 2024, while more than 70% of employers in these settings reported difficulty finding suitable candidates (ILO 2024; Group 2024). Across contexts, firms struggle to recruit and retain workers, even as large portions of the labor force remain underemployed with limited economic mobility. Employers often rely on referral-based hiring, where incumbent workers recommend potential candidates from their social networks, a practice that is particularly prevalent in developing countries given higher frictions (ILO 2018; Chandrasekhar, Morten, and Peter 2020; Breza and Kaur 2025). However, the reliance on referrals limits the scope for workers without connections to obtain work. This is especially relevant in contexts with spatial misallocation of job opportunities and where workers migrate to transition across sectors. In this setting, firms who rely on referrals may end up with a low-search cost yet high-turnover equilibrium, effectively barring minority workers who have a high willingness to accept these jobs. Such an imbalance points to a market failure: referral allocations skewed toward majority groups can trap firms in costly cycles of turnover and leave marginalized candidates on the sidelines.

This paper studies the tension between balancing the demographic composition and associated cohesion costs from diverse teams, and the gains from improving expected retention. When firms concentrate referral opportunities among frequent referrers, they may reduce marginal search costs but risk recruiting workers with strong outside options—leading to higher turnover (Calvo-Armengol and Jackson 2004; Bolte, Immorlica, and Jackson 2021).[1] Conversely, firms might worry that raising the share of underrepresented workers could trigger inter-group conflict and hamper coordination (Hjort 2014; Ghosh 2023; Lazear 2000). I examine these competing forces in entry-level employment in the context of India's caste system – an institution historically designed to restrict occupational mobility, which still influences many aspects of economic life (Cassan, Keniston, and Kleinberg 2022; Oh 2023; Lowe 2021; Munshi 2019).

This paper uses an experiment at a manufacturing firm in Delhi – and a complementary lab-in-field extension – to test whether redirecting referral opportunities toward lower caste workers can boost retention and productivity. At the firm, over 90% of frontline workers are hired via referrals from incumbent employees, yet lower-caste workers—who make up 21% of the workforce—receive only 12% of referrals. Ordinarily, team supervisors decide which incumbent gets to refer when a vacancy arises. During the experimental period, these referral opportunities were randomly assigned to only lower caste incumbents in treatment teams, whereas referrals were allocated randomly – without consideration of the incumbents' caste – in control teams.

Compliance with the listed referral allocations in each group was high: 92% on average on a team-by-month level. Over the six months of the experiment, lower caste representation in treatment teams rose by 15 percentage points, a 62% increase (p-value =

[1] Here, "network quality" refers to workers having more frequent or valuable links, typically to incumbents at desirable jobs.

0.01). Correspondingly, retention improved in treatment teams, with monthly churn – defined as the sum of exits and entries – falling by 4 percentage points (a 41% drop) compared to control teams. This reduction translates into 0.6 fewer exits and 0.7 fewer new hires over the period, driven primarily by newly hired lower caste workers staying longer. In contrast, control teams continued losing new hires at roughly their baseline rate. Survey data from the baseline reveal that lower caste referrals were 65% more likely to be unemployed and more likely to work in rural settings than their upper caste counterparts—highlighting the difference in these groups' outside options and suggesting an explanation for why lower caste entrants tend to remain at the firm for longer.

The effects of improved lower caste representation on productivity are ambiguous ex-ante. The first consideration is inherent ability: lower caste workers may have lower human capital in expectation, even as there have been improvements in recent decades in their educational attainment (Hnatkovska, Lahiri, and Paul 2013). The second potential drawback is that a more mixed team may be less productive owing to worse social cohesion. The literature on the effect of workplace diversity on output suggests that there are (at least short run) declines in mixed teams, largely stemming from additional coordination costs or risks of conflict (Hjort 2014; Ghosh 2023).[2] On the other hand, there are expected benefits from the reduction in turnover discussed above, and from potential complementarities. I find that the next effects are significantly positive in my setting.

Over six months, team-level productivity increased by 7% (0.07 σ, p = 0.03), with especially large gains in tasks requiring coordination among teammates. This result aligns with evidence that churn is most disruptive when workers depend heavily on one another (Hamilton, Nickerson, and Owan 2003a; Huckman and Staats 2013). Notably, I do not detect any short-run declines in output. In fact, output in treatment teams is similar to control teams initially and improves meaningfully over time, suggesting that the reduced turnover allowed teams to refine their collaboration and boost efficiency. I find no direct evidence of declines in survey measures of cohesion. At endline, neither job satisfaction nor social cohesion differed significantly between treatment and control teams.[3]

These findings raise two questions. First, given their incentives to improve productivity, why were supervisors not allocating referrals to lower caste workers already? In detailed surveys at baseline where I elicit their priors beliefs about the effects of lower caste entrants on cohesion, the supervisors report fears that reflect the findings from the

[2]Hjort (2014) uses short-run variation in the ethnic composition of teams in a manufacturing firm in Kenya and shows lower output in diverse teams. Ghosh (2023) experimentally varies the religious composition of teams in a manufacturing plant in West Bengal, India, and shows that among teams where workers must depend on each other more, there are noticeable short-run declines in output. In this setting, the longer-run effects of diversity are positive. I replicate the focus on the coordination level of tasks in my setting, stratifying treatment assignment on this dimension.

[3]While self-reported measures can have limitations, these results imply that any negative impact on worker cooperation from increased diversity was minimal or nonexistent in this setting. I address the shortcomings of self-reported outcomes through revealed measures of cohesion in the lab-in-field experiment.

literature on mixed teams. In particular, they worry that there will be costs to cohesion and output if they allocate referrals to lower caste incumbents.[4]

This leads to a second natural question: why are cohesion costs – common in the diversity literature – absent in my setting? I explore this question by testing worker perceptions of various recruitment policies through vignette surveys. Participants, drawn from a population similar to the firm's workforce, rated how adding team members would affect cohesion and productivity under a wide range of recruitment modes. The results show that outsider hires (lower caste recruits unlinked to current team members) are viewed as disruptive for cohesion and productivity, while when the same lower caste recruits are hired through referrals from incumbent workers,they are perceived as harmless to team cohesion. These perceptions mirror the firm experiment's findings: shifting referrals to underrepresented incumbents does not incur the cohesion losses often anticipated in more general diversity interventions.

To further validate these findings, I supplement the vignettes exercise with an incentivized lab-in-field experiment. Conducted with nearly 500 workers drawn from the same population as the firm's workforce across 12 workshops, this study setting provides tight control over task structure, data collection, and recruitment policies—thereby closely mirroring the firm's conditions. Each team starts with three workers, one lower caste and two upper caste, then completes two production rounds separated by a recruitment spell that brings the team size from three to four. I test three recruitment modes: (1) referral by the incumbent lower caste worker, (2) referral by an incumbent upper caste worker, and (3) direct hiring of a lower caste outsider. The first two mirror the treatment and control arms of the firm experiment, whereas the third replicates common diversity-focused recruitment policies that ignore incumbent links.

Results from the lab-in-field extension echo the firm experiment and clarify why cohesion does not decline. Lower caste referral teams and lower caste outsider teams start with the same production levels in round one, experience the same nominal increase in diversity, and recruit observationally similar new hires.[5] Yet labeling the new lower caste hire as a referral from an incumbent teammate vs. an unconnected outsider causes large productivity differences. In outsider teams, the fourth worker adds no statistically significant team-level output, whereas in referral teams, output rises by 22% ($p < 0.01$). As in the firm, these gains emerge only for high-coordination tasks, reinforcing the notion that networked integration avoids cohesion frictions and enhances performance.

Unlike in the firm setting, in the lab-in-field experiment I am able to deploy revealed-preference measures of cohesion, adapted from prior behavioral games (Schofield 2014; Breza, Kaur, and Shamdasani 2017). These measures capture how pairs of workers

[4] I am not able to rule out other reasons why supervisors do not allocate referrals to lower caste incumbents, including preferences against mixed teams that could overcome the potential perceived gains in output. Yet another mechanism could be risk aversion, wherein supervisors are disproportionately worried about short-run costs, even as the firm's owners may not be, reflecting a principal agent problem within the firm.

[5] Following Pallais and Sands (2016), this design holds selection constant, isolating the effect of recruitment mode.

(randomly matched across rounds) cooperate within each team. Lower caste referral teams exhibit 9% higher cohesion than lower caste outsider teams ($p = 0.02$), and this gap appears among both incumbent-incumbent pairs and pairs involving a new hire. Thus, the mode of recruitment shapes team-wide cohesion, not just the relationship between the referrer and the recruit. Finally, disaggregating by caste composition shows no effect in same-caste pairs but a strong positive effect for cross-caste pairs in lower caste referral teams.

The referral-based recruitment policy evaluated here has lasting effects on the firm. More than five months after mandatory referral allocations through the experiment ended, retention and productivity remain higher in treatment teams, and supervisors continue directing referral opportunities to lower-caste workers. This suggests that treatment supervisors learn from their exposure to the referral allocation policy mandated in the intervention. The share of referrals going to lower caste workers in these teams still exceeds their (now higher) fraction of the team's total workforce post-experiment. Conversely, control teams largely revert to pre-experiment allocations. I develop a dynamic model of referral allocations to study the evolution of supervisor beliefs. In this model, supervisors hold prior beliefs about cross-caste cohesion costs and may exhibit taste-based discrimination against lower castes. Calibrating these priors with baseline survey data and allowing Bayesian updating under alternative referral policies, the model shows that sufficiently high initial beliefs about cohesion costs can prevent supervisors from ever allocating referrals to lower caste workers – even when improving team stability would be more profitable in the long run.

This paper contributes to a large literature on labor market dynamics, social networks, and workplace diversity, with policy implications for contexts where market segmentation influences hiring. First, it builds on foundational work on job search models (Jovanovic 1979; Montgomery 1991), and the costs faced by firms in developing countries due to high turnover rates (Donovan, Lu, and Schoellman 2023; Tybout 2000; Skelton, Nattress, and Dwyer 2019; Shiferaw, Söderbom, and Zewdu 2020; Halvorsen 2021). It is also closely related to the body of work on the use of social networks in worker recruitment (Afridi, Dhillon, and Sharma 2021; Chandrasekhar, Morten, and Peter 2020; Osman et al. 2021), and extensive research comparing referred vs. non-referred workers – mostly from matched employer-employee data in developed countries (Kugler 2003; Dustmann et al. 2016; Burks et al. 2015; Brown, Setren, and Topa 2016) – and from developing-country studies on referrals as a device to mitigate adverse selection and moral hazard (Heath 2018; Pallais and Sands 2016), or explain how referrers choose candidates (Beaman and Magruder 2012; Beaman, Keleher, and Magruder 2018). I extend this literature by examining an understudied dimension: how referral opportunities themselves are allocated among incumbent workers, along with downstream effects on retention and productivity. This work is also related to another foundational strand in labor economics: spatial misallocation as a result of disadvantaged groups with potentially high willingness to accept jobs not being able to access them (Kain 1968; Chetty et al. 2014; Moretti 2012).

Methodologically, the experiment is unique in combining firm-level variation in referral practices. Previous research has largely relied on administrative or archival data –

where referral links are only inferred – or on stylized settings where researchers directly act as employers. Neither fully captures the employer–incumbent worker relationship that shapes both referral choices and new hires' integration. By working with an existing workforce to alter hiring policies, this study offers new insights into how referral allocations evolve and impact firm outcomes.

Second, this work speaks to a broad literature on cohesion and workplace productivity (Roethlisberger and Dickson 1939; Lazear 2000; Mas and Moretti 2009), including studies of managerial roles and recruitment policies (Fershtman and Gneezy 2001; Prat 2002; Bandiera, Rasul, and Barankay 2005; Bandiera, Barankay, and Rasul 2009). A parallel strand examines ethnic diversity and its impact on trust (Alesina and La Ferrara 2005; Ottaviano and Peri 2006) and firm productivity (Hjort 2014; Ghosh 2023). This work presents the potential of a recruitment strategy that blunts the adverse impacts of diversity on cohesion, while retaining the value that both firms and workers see in referrals.

In addition, this research connects to the sociology and social psychology literatures on "individuation" vs. group stereotypes and the contact hypothesis. That literature finds that close interactions with out-group members can lead observers to view them as unique exceptions rather than revising beliefs about the entire group (Fiske and Neuberg 1990; Granovetter 1973). Subsequent "subtyping" may then preserve negative stereotypes about the broader out-group (Brewer 1999). Here, I show that referral-based hiring can extend positive associations from one lower-caste worker to direct links of that worker,even without fully shifting generalized beliefs. This paper also has key theoretical antecedents in Calvo-Armengol and Jackson (2004) and Bolte, Immorlica, and Jackson (2021), incorporating their models of network-dependent labor markets into an experimental setting that tests predictions on referral allocations, cohesion, and productivity.

Lastly, this paper has a range of widely relevant policy implications. Across the developing world, firms heavily rely on referrals (Afridi, Dhillon, and Sharma 2021; Afridi, Dhillon, Xin Li, et al. 2020; Witte 2021). Network-based hiring tends to amplify homophily, which can perpetuate inequality and sustain inefficient equilibria with high turnover and underutilized talent pools—particularly for those excluded from incumbent networks.[6] Expanding referral opportunities to underrepresented groups can disrupt these inefficiencies by tapping workers with higher willingness to accept jobs, which feeds into higher expected retention.

There is the potential for large gains from the workers' perspective as well. Lower caste groups – those at the bottom of the hierarchical caste system – remain underrepresented in high-premium jobs, relative to their share of the overall population.[7]

[6]Parallel patterns appear worldwide. A broad literature links referral practices to unequal labor-market access, including gender disparities in developing countries (Beaman, Keleher, and Magruder 2018), racial gaps in the United States (Arrow and Borzekowski 2004), and immigrant exclusion (Munshi 2003; Patacchini and Zenou 2012).

[7]My analysis of the nationally representative India Human Development Survey (IHDS) 2014–15 finds that lower castes constitute more than 30% of agricultural laborers – an industry with the lowest

This gap suggests a segmented structural transformation in India, with caste continuing to drive occupational segregation (Munshi 2019). Historically limited access to higher-premium urban jobs for lower castes – compounded by referral practices where incumbent workers refer peers from the same caste – offers one explanation for why these groups remain largely confined to less-remunerative sectors. Disrupting this status quo equilibrium promises to improve the labor market outcomes for these groups.

The rest of the paper proceeds as follows. In section 2, I discuss the background for this context and the experimental setting, and some descriptive motivating evidence through stylized facts. In sections 3 and 4, I present the experimental design and the empirical strategy respectively. I present an overview of a dynamic theoretical model rationalizing the firm's referral allocation process in section 5. I discuss the results from the firm experiment and the lab-in-field extention in section 6. I conclude in section 7.

2 Background

2.1 Setting for the firm experiment

This experiment is conducted in a footwear manufacturing plant just outside of Delhi, in a region that has a high density of footwear, textiles, and other similar manufacturing plants. This industrial cluster is about 30 kilometers from Delhi's central business district, and primarily attracts migrant workers from neighboring states. These workers often reside in informal settlements in and around the industrual cluster, and tend to have previously worked in agriculture before migrating to the city for their current roles.

The firm employs between 700 and 900 workers on average, with the precise rate varying month-on-month based on seasonality in demand. Frontline workers who are actively engaged in the production process are called "helpers". Production takes place in teams, with team size typically varying between 5 and 8 helpers. Each team is led by a supervisor, who is only nominally engaged in production himself, and is resopnsible for recruitment, training, troubleshooting production challenges, and other miscellaneous tasks. Helpers tend to work on informal contracts, wherein they are paid a fixed monthly wage with no explicit financial incentives, and are able to leave the firm and occassionally re-join it. Supervisors are on permanent long-term contracts with the firm.

Each team works on a discrete task, with the tasks summarized below. Helpers and supervisors both generally stay at one team during their tenure at the firm. There are virtually no opportunities for helpers to get promoted to the level of supervisors. Similarly, supervisors are generally not given opportunities for promotion, except in rare situations.

wage residual across the entire Indian economy (conditional on location, age, gender, and education) – and fewer than 10% of manufacturing workers.

Task	Description	# of teams	Coordination level
Assembly	Preparing finished footwear pieces	27	High
Packing	Packing finished articles for shipping	22	High
Sole preparation	Layering the components that make up the sole of the footwear	20	High
Maintenance	Repairing machines across the factory	12	High
Cutting	Portioning fabric and other materials for subsequent steps	10	Low
Polymer	Preparing pellets for extruder and shaping into basic footwear structure	21	Low
Stitching	Stiching straps and logos onto footwear	21	Low
Total		132	

2.2 Referrals in developing country labor markets

In developing economies, informal mechanisms dominate recruitment processes, particularly for low-skilled and entry-level manufacturing jobs. Primary among them is the use of referrals from incumbent workers, which is particularly prevalent in settings where formal hiring processes are underdeveloped, and employers face significant informational asymmetry with respect to the reliability and productivity of potential workers.

Referrals are especially common in manufacturing sectors where tasks are repetitive, entry-level positions tend to be abundant, and worker misconduct or mistakes can severely disrupt productivity. In such settings, referrals fulfill the dual role of improving the firms' screening and mitigating moral hazard. Over 70% of jobs in the manufacturing sector in India are estimated to be sourced through referrals, with similar proportions in other contexts in South Asia and Sub-Saharan Africa (Witte 2021; Caria, Franklin, and Witte 2023; Beaman, Keleher, and Magruder 2018).

Employee referrals in these contexts are not merely a convenience; they serve as a vital informal institution to address key challenges faced by firms. For instance, in environments with weak legal enforcement and limited formal monitoring systems, firms often struggle to ensure worker discipline. Referrals leverage the social ties between referrers and referrals to create mechanisms of informal accountability. When a recruit underperforms or engages in misconduct, the referee may face social or reputational repercussions, providing a natural incentive for referees to recommend reliable candidates. This dynamic can reduce the need for costly monitoring systems and efficiency wages, making referrals particularly attractive to employers in resource-constrained settings. The reliance on referrals also reflects the socio-economic structure of many developing countries, where social networks play a central role in economic interactions.

The proportion of jobs secured through referrals varies across regions. For instance,

studies indicate that in the United States, approximately 13% of workers find employment via social contacts, while in the Philippines, this figure rises to 72%. In developing economies, the reliance on social networks for job searching is even more pronounced. The World Bank reports that the vast majority of jobs in these regions are obtained either through social networks of friends and relatives or by workers directly approaching employers to inquire about openings (Demirgüç-Kunt 2011).

Moreover, referrals not only serve the interests of firms but also offer substantial benefits to workers by reducing the costs and frictions associated with job search. For many job seekers – particularly those facing significant barriers such as limited access to formal recruitment channels – their social networks become the primary conduit to learning about employment opportunities. The frequency and quality of job leads that workers receive depend critically on these networks. In this system, it is plausible that individuals with higher quality social networks would hear about higher quality job openings, affording them better outside options. In contrast, workers with weaker networks and fewer alternatives would be left with fewer prospects, perpetuating a cycle where only those already advantaged gain access to the most desirable positions. Thus, the mechanism that allows firms to lower their marginal search costs – by leveraging referrals to target workers with dense connections – may simultaneously reinforce labor market disparities.

2.3 Caste

The Indian caste system is at least 3000 years old. The delineation of society into four hierarchical categories ("varnas") - *Brahmins*, *Kshatriyas*, *Vaishyas*, and *Shudras* - is found in the book *Manusmriti*. The *Dalits* (previously considered "untouchable" by the traditional upper caste groups) exist outside of this hierarchy altogether, and are the lowest social group. Each *varna* in turn has hundreds of sub-groups known as "jatis", such that jatis within varnas are also categorized hierarchically. In aggregate, there are estimated to be upwards of 4000 jatis in India (Deshpande and Ramachandran 2024). These jatis tend to be characterized by strict rules of endogamy, with genetic evidence establishing that strict endogamy has been practiced for thousands of years, helping sustain the institution of caste (Moorjani et al. 2013). The caste system also enforces norms around social interactions, such as restrictions on sharing food, as well as residential segregation and most importantly from the perspective of this paper, strict boundaries around occupational roles (Oh 2023; Lowe 2021).

While the *jatis* form the basis of the caste system, modern statutory government classifications—General, Other Backward Classes (OBC), and Scheduled Castes (SC)—are significant in studying caste-based discrimination in India (Munshi 2019). These classifications align with the aforementioned hierarchy. Broadly speaking, General castes rank above OBCs, while OBCs rank above SCs. Another group of importance is indigenous Indian communities, categorized as the Scheduled Tribes (ST). While there are noteworthy differences in the historical status of these groups, they are ranked in policy and inter-group interactions approximately at the same level as SCs. In this paper, I

will pool together SCs and STs into a "lower caste" category, while the reamining are classified into the "upper caste" category.

SCs continue to face systemic discrimination in India despite constitutional protections. The caste system's rigid social hierarchy marginalizes SCs in several spheres, including access to education, employment, and social interactions. Practices of untouchability, though outlawed, persist in many rural and urban areas. For instance, 24% of households across India still practice untouchability, a custom where physical contact with SCs is avoided to maintain "purity." This form of exclusion restricts SCs' access to shared community resources and perpetuates social isolation (Desai and Vanneman 2011). Moreover, these caste-based differences are intertwined with spatial segregation. SCs often reside in separate hamlets within villages, and inter-caste interactions are minimal, leading to limited access to social capital and networks (Anderson 2011).

Economic opportunities for SCs are often limited due to occupational stigmas and discrimination in hiring. While there has been some improvement in the economic standing of SCs over recent decades, owing in part to affirmative action programs in education and public sector employment, disparities remain significant(Hnatkovska, Lahiri, and Paul 2013; Asher, Novosad, and Rafkin 2024). Furthermore, SCs are frequently relegated to the most physically demanding and least desirable jobs, reinforcing their historical economic subjugation. Approximately 50% of economic exchanges, such as food transfers and loans, occur within the same caste (jati), highlighting the role of caste in structuring social and economic life (Munshi and Rosenzweig 2015). This caste-based homophily results in SCs being underrepresented in higher-paying sectors. As a result, SCs often remain confined to low-paying, menial jobs with limited upward mobility, perpetuating economic disparities and social marginalization. In independent analysis of nationally representative data from the Indian Human Development Survey of 2011-12, I find that lower castes are disproportionately represented in landless agricultural labor and in rural areas more generally (figure B6). Even when they are able to migrate to cities, they are concentrated in low-premium industries.

2.4 Stylized facts

In this section, I highlight descriptive evidence on occupational segregation, the role of networks in job search, and caste dynamics at the site(s) of the experiment, drawing primarily on the detailed baseline survey conducted at the firm, and supplementary analysis from other data sources.

Fact 1: LCs are underrepresented in manufacturing work generally, and at higher levels of the firm's workforce in particular. As referenced above, I use the 2014-15 round of the IHDS to analyze the representation of LCs across industries. I find that controlling for education, district, age, and gender, LCs are more likely to be found in landless rural agricultural labor. Within urban areas, they are disproportionately more likely to be found in industries with low premia (figure B6). In the firm where the experiment is conducted, LCs make up 21% of the overall workforce at the firm, but

the proportion is only 6% at the team leader level, and none at the level of the firm's executive.

Fact 2: Referrals play an outsize role in matching workers with jobs in entry-level jobs in developing country contexts. Referrals are the primary mode of hiring in the firm where the experiment is conducted, with 92% of the helpers (i.e. the workers actually engaged in production) at the time of the baseline survey having been hired through referrals, while the rate is approximately 81% when team leaders and upper level executives are also considered. This is not atypical: Afridi, Dhillon, and Sharma 2021 show that 75% of the workers in their setting (garment manufacturing plants in India) find their jobs through referrals; Sapin et al. 2020 report an estimate of 72%; and Kumar and Srivastava 2025 find rates of 71%. Importantly, these estimates are all considerably higher than those from developed country settings. Schmutte 2015 report a rate of 10% in the US context. [8]

Fact 3: Referral opportunities are concentrated among upper caste workers in status quo. In the baseline survey at the firm, I ask the universe of all helpers at the firm the frequency of receiving referral opportunities. The probability of receiving a referral is more than thrice as large for UCs than LCs, while there is no detectable difference by the extent of education or experience at the firm (figure B7).[9] Controlling for education and experience closes the caste gap in referral access somewhat, though not meaningfully. Importantly, conditional on getting the opportunity to refer someone, a plurality of workers, irrespective of caste, refer their relatives to join the firm (48%), and a smaller share refer previous coworkers (22%) and other acquaintances (30%).

Fact 4: There is substantial clustering in the firm's workforce in status quo, and worker networks are highly saturated by caste. As of the baseline, the firm's 724 workers as of baseline are clustered into just 137 unique referrer-referral links. The average worker has 1.64 links across the firm's workforce (this is equivalent to the average degree centrality of the network), and 1.05 of these are made up of the same caste as the nodal worker. The density of the network is 0.0021, indicating relatively low levels of aggregated connectedness across the graph. A network graph representing these referrer-referral links is in figure B1, and figure B2 presents a benchmark for what the network would look like if each referrer was equally likely to refer someone from either caste, while keeping the total number of UCs and LCs constant. This homophily in the workforce extends to the links of potential referral candidates, wherein most incumbent workers claim to share prospective referral links with others from their own

[8] These estimates are primarily from worker surveys *across* firms. In my setting, the share of workers matched through referrals is higher because it is concentrated within a specific factory. In general, recruitment policies are likely to vary much more across firms than within, which could explain the somewhat lower rate of referral dependence in a worker survey.

[9] Education by itself does not vary meaningfully among the worker population, since the firm requires all its workers to have studied till at least the eigth grade. Only 11% of the workers at baseline report studied less than this requirement. Note that I am unable to identify instances of misreporting in the status of education, but anecdotally the firm itself also does maintain a rigorous standard of evidence that the workers must furnish.

caste (figure B3).

Fact 5: LCs have a higher retention rate at the firm in status quo, owing primarily to worse outisde options. At baseline, I find that incumbent LCs are likely to have longer survival and retention rates at the firm (Kaplan-Meier survival $\chi^2 = 34$, p-value < 0.01, spending 30% longer at the firm on average; figure B4). These patterns are also reflected in the previous jobs reported in the cross-section of the firm's respondents at baseline. They are 20% (p-value = 0.04) less likely to be in fixed monthly employment, and 81% (p-value = 0.049) more likely to be in daily wage labor. Their aggregate earnings in previous jobs were also lower (by 12%, p-value = 0.092).

Fact 6: LC and UC incumbent workers have similar extensive margin referral preferences, but the characteristics of their potential referrals vary markedly. As part of the baseline, I also collected detailed network information on potential referral candidates, asking respondents to list the names and some basic information of an (unconstrained) set of their contacts that they would bring to the firm if they ever received a referral opportunity. Both UC and LC respondents list the same number of contacts (2.02 on average), are highly likely to want to refer someone from their own caste group (86.5% on average), and are highly likely to want to refer someone from their close or extended family (83.1% on average).[10] However, LCs are likely to report having had 5 more conversations with their referral candidates in the previous 30 days than UCs (p-value = 0.097), likely reflecting stronger links. Importantly, LC workers' referral candidates are substantially less likely to be in urban work (13%, p-value = 0.006), and are less than half as likely to already be in Delhi (p-value < 0.001). This underlines the fact that most LCs are picking contacts from those who are predominantly in rural areas, while UCs are generally considering those already in urban job-seeking cycles. Despite these differences, both UCs and LCs rate the inherent ability of their referral candidates similarly.

Fact 7: Supervisors worry that LC entrants would diminish team cohesion, but these views are not shared by the workers themselves. In baseline surveys with supervisors, I present them profiles of three hypothetical workers presented in random order, such that two of them have the same education, experience and state of origin, but differ in their names' caste labels. For each of the hypothetical profiles, supervisors are asked how they expect them to perform on the team on a range of attributes: their net productivity, how hardworking they would be, whether they would be easy to train, and whether they would get along with the team. The latter is meant to be a measure of the effect of the new entrant on the team's cohesion. I find that while supervisors do not believe that LCs would be differentially hardworking or easy to train, they would be less likely to get along with their team (by 5%, p-value = 0.014). This results in a net lower expectation of LCs' productivity (by 3%, p-value = 0.029). All of these measures were also elicited in the baseline survey with incumbent workers, which I use to generate the workers' beleifs about each caste group on productivity, hardwork, and cohesion. I do not find any statistically significant differences on any dimension,

[10] These indicators are not statistically different for UC and LC respondents.

indicating that workers do not share their supervisors' apprehension that LCs would lower cohesion.

I also attempt to recover these insights from supervisors on their actual workers, since there may be other unintended biases in their consideration of hypothetical workers in the stylized exercise. On subjective ratings for their current workers across a range of indicators, I find precise nulls on the caste comparison, suggesting that they view all of their workers similarly across these dimensions. This is consistent with positive selection (i.e. LCs who are actually on the team have higher ability than hypothetical LC workers who will be drawn from the broader distribution), training (i.e. they believe the LCs are now trained and better than a draw from the distribution), and/or contact theory (i.e. their priors about their specific workers have improved as a result of extended exposure) (Lowe 2021). In a separate exercise as part of the baseline, I share with supervisors profiles of three randomly selected workers currently on their team, and ask them to rate potential referral candidates from each of these workers on a range of dimensions. This exercise suggests that they believe their LC workers would refer others who are similarly qualified, but would be more likely to accept the job offer (6%, p-value = 0.018) and would be less likely to get along with the rest of their team (8%, p-value = 0.027).

Fact 8: Out-group biases. At baseline, there is no clear correlation between reported cohesion and worker or team characteristics (such as worker caste, experience, and team size or caste composition). However, LC workers report having more friends in teams with below median LCs (92%, p-value = 0.059), indicating some selection and differential group dynamics in teams with fewer LCs. Additionally, there are no systematic biases in workers' or supervisors' beliefs about LCs, elicited broadly through the discussion of hypothetical profiles that are similar on observable dimensions like education and experience, but differ by caste. The profiles are presented in figure A9.

3 Experimental design

I explore the effect of referral-based hiring systems on team cohesion and productivity using two experiments. The primary experiment is conducted at a manufacturing firm in an industrial cluster close to the city of New Delhi in India. This firm is engaged in the production of footwear, and typically employs between 700 and 900 workers in the actual frontline production process (called "helpers"). Production is organized in terms of teams, each involving between 5 and 8 workers. Each team works on a discrete task, and tasks vary in the extent of coordination that works must exert in order to perform them. Workers are paid a fixed monthly wage, and have no performance-linked bonuses. Each team is led by a supervisor, who earns a higher wage and is responsible for streamlining production and personnel issues for the team. In the pre-experiment status quo, 92% of the workers at the "helper" level are hired through referrals. Referral opportunities are allocated by supervisors to incumbent workers per their discretion.

This experiment is implemented at the team (and thus supervisor) level. For the duration of the experiment, which was announced to supervisors after the baseline surveys, I took away the discretionary powers of *all* supervisors to allocate referrals. In treatment teams, each time there was a vacancy in the team either due to an exit or increased labor demand, referral opportunities were allocated to a randomly chosen incumbent lower caste worker. In control teams, referral opportunities were allocated to a randomly chosen incumbent worker irrespective of their caste.[11] This design is represented in figure A2.

In the baseline, I attempted to survey the universe of all workers and supervisors at the firm, successfully managing to reach 94% and 92% respectively. The worker baseline surveys included a detailed network elicitation exercise, wherein respondents were asked to list the names and some basic demographic information for contacts that they would bring to the firm if they received a referral opportunity. Thus, prior to the allocation of teams to one of the two treatment groups, I was able to construct a network of potential referrals for the near-universe of incumbent workers.

The experiment lasted six months, wherein referrals were allocated based on each team's treatment status. Table 2 reports results on the compliance with randomization, with respect to the caste status of the recipients of the referral opportunity. Compliance with listed referral allocations was high on average, across both treatment and control teams.

A timeline of each implementation and survey component of the study is listed in figure A1. Following the intervention at the firm, I conducted a detailed endline survey with the universe of workers and supervisors, reaching 92% and 93% respectively. The primary outcomes from the firm survey, however, come from administrative data. This includes weekly team-level data on the units produced or processed by each team. In addition, I use personnel data at the team-by-month level to generate measures of the teams' caste composition, tracking entry and exit by incumbent status as well as caste group, and the links between referrers and referrals.

Following the conclusion of the experiment at the firm, I implemented a lab-in-field experiment (figure A4). This experiment was conducted with workers drawn from a population similar to the workers at the firm. The lab-in-field experiment was conducted over 12 spells, with each spell including two 45-minute production sessions and surveys before and after the production sessions. Each spell of the experiment starts with 30 workers, 10 of whom are lower caste and the remaning are upper caste. These workers are sorted into 10 teams, each with one lower caste and two upper caste workers. Each worker is asked for a potential referral candidate, i.e. we ask them to connect us to someone who might be asked to work with them in the next session. Following this elicitation, workers perform a task in teams. The task was designed

[11]The implementation of this referral allocation, across both treatment and control teams, was done in a centralized way by the firm's HR department. At the beginning of each month, I shared an ordered list of which incumbent worker each team's supervisor should call upon in the event of a vacancy that needs to be replaced. An example of this ordered list is in figure A3.

to vary the extent of coordination across teams.[12] Each worker was given fabric of a certain color, and the team was asked to cut strips of this fabric of pre-determined dimensions, and then tie these strips together to make a rope as long as they could in the 45 minutes. In high-coordination teams, workers were asked to follow a particular order of colors, while the order was immaterial in low-coordination teams. Workers' payoffs were linked to the overall performance of the team, i.e. they were given a fixed payoff of INR 600 (USD 7) for participating in the lab-in-field experiment, and INR 5 (USD 0.06) per strip for the rope they are able to construct as a team.

Each production session was followed by a round of surveys with workers, where I elicited measures of social cohesion. Following the first production session, teams were randomized into three groups, wherein the groups differed in the mode in which a fourth worker was added to the team. In the first group, "LC referrals", the fourth worker is the referral of the incumbent lower caste worker. In the second group, "UC referrals", the fourth worker is the referral of one of the two upper caste workers, chosen randomly. In both the first and second groups, we communicate an identical message to the incumbent workers, i.e. a fourth worker is being recruited by inviting a referral from one of the three incumbent workers. In the third group, "LC outsiders", the fourth worker is not connected to any of the incumbent workers, but is thte referral of an incumbent lower caste worker from a *different* team. This mode of recruitment is communicated to the incumbent workers as a fourth worker being added to the team in order to balance the team's caste composition. The first two groups are designed to replicate the firm experiment, wherein "LC referrals" and "UC referrals" map on to the treatment and control teams from the firm respectively. The "LC outsiders" group is designed to replicate a non-referrals based diversity intervention, similar to common affirmative action programs and other policies that make the identity of the minority group salient, found both in the literature and in firms across the world. Allocation to these treatment conditions was stratified by the high vs. low-coordination nature of the task.

At the end of the second production session, workers are asked to rate cohesion and satisfaction. In addition, workers are organized in pairs and play a cooperation game in order to arrive at a revealed measure of their cohesion, modeled on Breza, Kaur, and Shamdasani (2017). Each worker is paired with another randomly selected worker from their team over two rounds, and they are asked to perform a task wherein they must work together to match the number of symbols that are in common across their two sheets.

The key comparison of interest in the lab-in-field experiment is "LC referrals" vs. "LC outsiders". Note that this design entirely switches off the effect of selection on productivity and cohesion, similar to Pallais and Sands (2016). All referral elicitations are made prior to the allocation of treatment, and the new workers in the "LC outiders" group are also referrals of similar incumbent LCs as in the "LC referrals" group. The only salient difference across the groups is the mode of hiring, and how this is communicated to the incumbent workers.

[12] A pictorial representation of the tasks is in figures A6 and A7.

4 Emprical strategy

In analyzing data from the firm experiment, I follow a simple regression specification:

$$Y_{imt} = \alpha + \beta \cdot T_m + \gamma_s + \phi_t + \varepsilon_{imt} \tag{1}$$

where Y is an outcome for an individual i from team m in time period t, T denotes a binary for the treatment status, γ denotes fixed effects for the randomization strata s, and ϕ denotes fixed effects for the time period. For outcomes that are defined at the team level, the i subscript is omitted accordingly. Standard errors are always clustered at the team level.

In the lab-in-field experiment, I employ a similar regression specification as in (1) for all survey-based outcomes and in the analysis of the revealed measures of team cohesion:

$$Y_{im} = \alpha + \beta \cdot T_m + Y^o_{im} + \kappa_t + \varepsilon_{im} \tag{2}$$

where indices are defined as in (1), T denotes a categorical variable for the three treatment groups, κ_t refers to fixed effects for the work spell, and Y^o_{im} refers to the value of the outcome collected in the screening or baseline survey prior to the production sessions, where available.

For analyzing the effects on team output in the lab-in-field extension, I employ the following specification:

$$\Delta_s = \alpha + \beta \cdot Z_s + \gamma_s + \kappa_t + \varepsilon_s \tag{3}$$

where $\Delta_s = \text{Output}_{s,2} - \text{Output}_{s,1}$ is the difference between the output at the team level between the first and second production sessions, Z denotes a categorical variable with three values, one for each type of team, γ denotes fixed effects for the randomization strata, and κ_t refers to fixed effects for the work spell,. Standard errors are always clustered at the team level. Note that I present the results on both the team-level output as well as the output per worker, where the latter is simply the former divided by three and four for the first and second production session respectively.

As a part of the lab-in-field experiment, I also collect data on a continuous measure of cooperation between workers on each team through a symbol-matching game. The effect on cooperation is analyzed with the following specification:

$$Y_{ipm} = \alpha + \beta \cdot T_m + \kappa_t + \varepsilon_{im} \tag{4}$$

where Y is the outcome for an individual i belonging to pair p in team m, T denotes a categorical variable for the three treatment groups, and κ_t refers to fixed effects for the work spell.

Lastly, I employ a similar specification for the endline follow-up survey with the referral

candidates elicited at the baseline survey at the firm:

$$Y_{jmt} = \alpha + \beta \cdot T_m + \gamma_s + \phi_t + \varepsilon_{jmt} \tag{5}$$

where j refers to a candidate listed by an incumbent worker at endline as a potential referral candidate, T_m refers to the treatment status of team m, which is the team that the referrer of j at the firm is in, and the remaining variables are defined as in (1).

Across all the specifications listed above, the standard errors are always clustered at the team level, which is the unit of treatment assignment.

5 Theoretical framework

In this section, I describe a dynamic model for a firm that must allocate referrals to incumbent workers to fill vacancies, balancing the value of retention and cohesion as they interact with the firm's output and costs. The goal of this model is to rationalize the pre-experiment status quo where the firm allocates referral opportunities disproportionately to UC workers, by making precise the trade-off involved between retention and cohesion under different referral allocation policies. A fully fleshed-out version of the model with proofs is detailed in Appendix D.

In each period, (a representative team at) the firm must allocate referral opportunities to either LC or UC workers when vacanies arise. n_l^t and n_u^t denote the number of LC and UC workers, such that the total size of the team is conserved: $L = n_l^t + n_u^t$.[13] The firm's producion function is a basic Cobb-Douglas, augmented to incorporate cohesion costs (Bloom, Sadun, and Van Reenen 2016; Bartelsman, Haltiwanger, and Scarpetta 2013; Hamilton, Nickerson, and Owan 2003b):

$$Y^t = [A \cdot C^t] \cdot (L^t)^\alpha \cdot (K^t)^\beta, \tag{6}$$

where C^t is the net cohesion cost in period t, defined as follows:

$$\begin{aligned} C^t &= 1 - \hat{\kappa}_t \cdot \rho_t - \hat{\tau}_t \cdot \Delta_t, \text{ where} \\ \rho_t &= \frac{n_l^t}{L} \cdot \frac{n_u^t}{L} = \lambda_t \cdot (1 - \lambda_t) \text{ and} \\ \Delta_t &= \mathbb{I}\{\lambda_t > \lambda_0\} \cdot (\lambda_t - \lambda_0). \end{aligned}$$

λ_t is the share of LC workers in the team in period t, λ_0 is the initial LC share of the

[13] The team size does not change over time or with composition, based on the empirical evidence from the firm experiment.

team, $\hat{\kappa_t} > 0$ is the cost of cross-caste cohesion, and $\tau > 0$ is the extent of (taste-based) discrimination against LCs. Note that cross-caste cohesion has two components: the product of the caste-specific shares of the team, and an indicator term that accounts for whether there are more lower caste workers in period t than there were at the initial stage, scaled by the difference in the proportion of lower caste workers in period t and the initial stage.

The firm's profit in each period is:

$$\pi^t = Y^t - w \cdot L^t - \mu \cdot X^t, \tag{7}$$

where w is a uniform wage paid to all workers, and μ is the cost of training each of the X^t new entrants in the team in period t. Note that the training cost μ captures both explicit resources dedicated to training new workers, as well as implicit costs associated with reduced temporary output in the team as workers adjust to new team-mates.

In each period, after allocating referral opportunities, the firm admits new workers such that referrers and referrals always belong to the same caste group.[14] After the new worker(s) join, the firm observes the actual cohesion in the team:

$$C^t = 1 - \kappa_t \cdot \rho_t - \tau_t \cdot \Delta_t + \epsilon_t,$$

where $\epsilon_t \sim N(0, \sigma^2)$.

Thus, in each period, the firm's beliefs on κ and τ evolve using standard Kalman filtering, assuming that the firm is Bayesian.

In each period, each incumbent worker has a probability of retention that is distributed as follows: $p_i^g \sim \text{Beta}(\eta^g, \phi^g)$, with the parameters of the distribution calibrated against the baseline mean and variance of the retention probabilities for each group $g \in \{l, u\}$. These retention probabilities in turn drive the exits in each team: $Q_i^{g,t} \sim \text{Bernoulli}(1 - p_i^g)$.

The Bellman equation for the firm's dynamic problem is:

$$V(S_t) = \max_{X_l^t, X_u^t} \left[[A \cdot C^t] \cdot (L^t)^\alpha \cdot (K^t)^\beta - w \cdot L^t - \mu \cdot Q^t + \delta \cdot \mathbb{E}[V(S_{t+1})] \right].$$

Constraints

[14] This assumption matches empirical findings, where referrals are shared with members of the same caste group 88% of the time.

1. Vacancy constraint, i.e. new referrals must equal the total number of exits:

$$X_l^t + X_u^t = Q^t.$$

2. Team size is conserved:

$$n_l^{t+1} + n_u^{t+1} = n_l^t + n_u^t + X^t - Q^t.$$

3. Cohesion cost is defined as follows:

$$C^t = 1 - \kappa_t \cdot \rho_t - \tau_t \cdot \Delta_t + \epsilon_t,$$

where the perceived cohesion cost C^t used in the decision is:

$$\hat{C}^t = 1 - \nu_{\kappa,t} \cdot \rho_t - \nu_{\tau,t} \cdot \Delta_t$$

Starting with initial moments on the firm's priors that are taken from baseline surveys, I numerically estimate the dynamic model to trace the evolution of their beliefs under different referral allocation protocols. The evolution of beliefs around κ and τ are represented in Appendix D. This exercise establishes that even when returns to team stability are significant, if the firm's priors around cohesion costs are large, they may continue to be insufficiently experimental in their referral allocations to LCs.

6 Results

Firm experiment

Workforce composition and referral allocations. The intervention significantly altered the composition of the recipients of referral opportunities, and the extent of workforce diversity in treatment teams. The probability of a referral opportunity going to an LC worker increased significantly from 24% to 92% in treatment teams. The latter indicates that over the course of the intevention, 92% of the referral opportunities in treatment teams were allocated to incumbent LC workers.[15] The number of LC hires through referrals in treatment teams increased by a corresponding 83%, rising from an average of 0.29 per month in control teams to 0.53 per month in treatment teams (table 2). As a result, by the end of the intervention, the share of LC workers in treatment teams was higher by 15 percentage points relative to control teams, increasing from 25% to 40%.[16]. Importantly, these changes in the composition are not associated with

[15]The remaining 8% corresponds to cases of imperfect compliance, which could be driven by adminstrative errors at the firm's HR department, an incumbent LC's refusal to accept a referral opportunity, and/or a supervisor overruling the prescribed referral allocation rule. Unfortunately, I am unable to distinguish between these three.

[16]By the end of the intervention, 28 treatment teams were composed of majority LC workers, relative to 3 in control teams

any significant differences in team size. These effects indicate that the intervention was largely implemented successfully, and substantially changed the workforce composition in treatment teams.

Supervisors continued to allocate referrals to lower caste workers in treatment teams, even after the conclusion of the intervention. In figure 2, I report the evolution of the share of referrals given to lower caste workers and the team's caste composition during and beyond the intervention. Referrals are almost always targeted at LCs in treatment teams during the intervention period as listed above, while they approximately track the share of LCs in control teams. This is because in the latter, the referral allocations are done at random without consideration of caste. It is worth noting that the control group represents a recruitment regime that has more LC referrals than what we might expect in the status quo, thus suggesting that the results are an underestimate of what we might expect relative to the status quo. As the share of referrals to LCs in treatment teams remains close to 1, the share of LCs in the team increases steadily. After the intervention ends, and the supervisors regain the ability to allocate referrals to incumbent workers discretionarily, treatment team supervisors continue to allocate referral opportunities to LC incumbents. While this could be mechanical to some extent, given the increased share of LCs in these teams, their allocations are at higher rates than the share of LCs. This is indicative evidence that supervisors learn about the true cohesion costs in this setting, and the benefits of retention from increased LC representation. Meanwhile, in control teams, LC referrals drop to rates below the share of LC workers in these teams.[17]

Retention. Consonant with the changes in the composition of treatment teams over the course of the intervention, I find a large treatment effect in improved retention in treatment teams, driven primarily by new hires. Treatment teams experienced a statistically significant 41% lower monthly churn, where churn is defined as the sum of exits and entries. Control teams lose 1.4 workers per month over the course of the experiment, and 0.8 of these are among new hires, i.e. workers hired any time after the baseline. Aggregate exit rates were 35% lower in treatment teams, a large and statistically significant difference. As noted above, this difference is most stark among new hires, where the treatment effect is 57% (table 4).

As a consequence of this improved retention, workers in treatment teams had longer tenures on average, with a statistically significant increase of 1.4 months (24%) compared to the control mean. Naturally, this effect on tenure was pronounced among LC workers, who were newer in treatment teams than in control teams (table B8).

Why are LC new hires more likely to be retained by the firm than UC new hires? I find that this is explained by the relatively large and systematic difference in outside options across the caste groups. In the baseline network elicitation, wherein I ask incumbent workers to list potential referral candidates, these differences in outside options are

[17] I cannot rule out the absence of learning in control teams as a whole. Note that the incidence of LC referrals in control teams is also random, since the likelihood of an LC incumbent worker receiving a referral opportunity is equal to the LC share of the population in expectation. While the aggregate LC share of referrals in control teams is lower in the post-intervention period, but this masks considerable heterogeneity with respect to the incidence of exposure to LC referrals during the intervention.

already apparent. I show descriptive comparisons between upper and lower caste incumbent workers' referral candidates on four key dimensions in Figure 3. Through incumbent workers' responses about their prospective referral candidates, I show that lower caste incumbent workers' referral candidates are 14% more likely to accept a hypothetical job offer, about half as likely to already be in Delhi, about 25% less likely to be employed in urban work anywhere, and about 65% more likely to be unemployed.

Output. Improved retention and more stable teams in the treatment group should, ceterus paribus, translate to higher productivity. However, a counterveiling factor is increased diversity: recent literature suggests that mixed teams are associated with (at least) short run productivity costs owing to higher cohesion costs (Ghosh 2023; Hjort 2014). Thus, it is not obvious which direction the treatment effects on productivity, if any, would materialize in. Importantly, there are two additional considerations: whether treatment effects on productivity evolve dynamically over time, and whether the effects are differential in teams that work on tasks that require a high or low degree of coordination among the workers.

I find that the intervention produced significant gains in team output, particularly in high-coordination tasks. Treatment teams exhibited a statistically significant 0.07 σ increase in output over the duration of the experiment, translating to a 3.6 percentage point improvement in performance. High-coordination tasks showed a stronger effect of 0.10 σ compared to 0.02 σ in low-coordination tasks, the latter not statistically significant at conventional levels (table 3). Importantly, the gains in output continue beyond the intervention. Pooled over the time period including the intervention and five months of the post-intervention period, the treatment effect on output rises to a statistically significant 0.10 σ, again concentrated among the high coordination teams.

Contrary to the existing literature on diverse teams' short-run productivity losses, there was no evidence of initial declines following the intervention. Positive treatment effects on the productivity emerged by week six of the intervention and strengthened over time (figure 1). I use the timing of the first exit faced by the team in the intervention period, which can reasonably be assumed to be quasi-random, to test the effect of retention on output. In figure B8, I show that improvments in team-level output intensified with more stability in teams. Together, these results suggest that the value of retention shows up in output, especially in teams where workers must coordinate. As control teams in high-coordination tasks keep losing workers as in the status quo, the dynamic treatment effect intensifies. This insight is consistent with a model of production where each time a worker exits a team with an established coordinated routine, even the workers who remain at the firm must learn to re-establish these routines with the new entrants. Turnover is disruptive in these settings, and gains from stability continue to accrue over time relative to control teams who keep losing workers at rates similar to the pre-experimental average.

Another notable feature of these dynamic results is the absence of any statistically significant downturn in output as a result of diversity. In order to understand the absence of short-run negative costs in this setting, it is important to understand the effects on team cohesion.

Team cohesion. At the endline survey, I collected a range of self-reported measures of social cohesion, summarized in table 9. In columns 1-2, I report treatment effects on a self-reported rating of team cohesion, i.e. where respondents rate how cohesive and cooperative their team feels to them on a scale of 0 - 10. Self-reported team cohesion is slightly higher (2%) in treatment teams, but the positive effect is largely concentrated among LCs. Next, I report treatment effects on the number of friends they report having at work. This is an unconditional response to questions where I ask them to report the number of coworkers they would spend time with outside work, and the number of coworkers they would spend time with at work (during lunch breaks, etc.). I do not find aggregate treatment effects on either dimension. However, LCs in treatment teams are statistically significantly more likely to report having friends at the firm that they spend time with outside (40%) and at the firm (26%). Relatedly, I find no treatment effects on job satisfaction, workers' expectation of leaving the firm, or job search activity (table 10).

Importantly, I elicit beliefs about out-groups with a random subset of workers at endline. These questions are administered by sharing profiles of hypothetical workers. These profiles are presented in figure A9, where each respondent hears about three hypothetical workers, two of whom have identical education and experience but vary in their names and thus caste status.[18] Workers are then asked to rate how productive these hypothetical workers might be, and how well they expect the hypothetical workers to get along with the team. I find no significent treatment effects on either dimension (table B7).

Overall, the firm experiment shows that directing referrals towards incumbent LC workers at the firm draws more LC hires in, improves retention in treatment teams, and manifests in higher output in these teams. Contrary to recent literature on diversity and the costs of cohesion in mixed teams, I find no noticeable negative effects on social cohesion in this team. One potential concern with these results is that the measures of social cohesion from the surveys with workers are all self reported and may suffer from measurement error and social desirability bias. Owing to logistical constraints, I was unable to administer more detailed behavioral modules measuring cohesion. I address this gap in the lab-in-field experiment, results from which are discussed in the following section.

Given the gains from this referral policy, a natural question one might ask is why the firm was not implementing this in status quo. One answer to this lies in the baseline surveys with supervisors, wherein I shared profiles of hypothetical workers of both castes and asked a range of questions eliciting supervisors' prior beliefs about these workers on a range of dimensions, including their perceived ability, how easy they would be to train, how long they are expected to last in the firm, and how well they might get along with the rest of the team.

[18] In this context, names are obvious signifiers of caste.

Vignettes experiment

How workers perceive the fairness and suitability of recruitment policies is an important determinant in their productivity. In contexts with different social group identities that lead to heterogenous teams, and where workers must coordinate to work on production, their beliefs about cohesion and by extension, productivity, are a useful indicator of the relative costs and benefits of different recruitment policies. In order to aid in developing a descriptive sense of this comparison, I conducted a vignettes experiment with 150 manufacturing workers, drawn from a population similar to the one that participated in the lab-in-field extension. These surveys involved sharing with respondents the profile of a hypothetical team engaged in a coordinated production task, modeled on the task from the lab-in-field extension. The respondents were told about the initial caste composition (one LC and two UC workers[19]), output at the individual and team level, and the initial self-reported cohesion. The information reported to the respondent is presented in figure A8. The respondent was then told that a fourth worker would be added to the firm, and they would be asked to estimate the output and cohesion in the subsequent production round. Each respondent was asked four of six potential recruitment policies in random order. The six recruitment policies presented to respondents are listed below, with the text in italics listing the actual content that was communicated to the respondents:

- LC referral (no caste label): *The new worker is hired through a referral given to [name of LC incumbent]*
- LC outsider (no caste label): *The new worker is unknown to existing workers, and his name is [name of new LC worker]*
- UC referral (no caste label): *The new worker is hired through a referral given to [name of a UC incumbent]*
- UC outsider (no caste label): *The new worker is unknown to existing workers, and his name is [name of new UC worker]*
- LC referral (caste label): *The new worker is hired through a referral given to the lower caste worker: [name of LC incumbent]*
- LC outsider (caste label): *The new worker is unknown to existing workers, is hired to balance the team by caste, and his name is [name of new LC worker]*

The set of recruitment policies in consideration is designed to cover the combination of three aspects: the identity of the entrant (UC vs. LC), the mode of recruitment (outsider vs. referral), and whether the caste identity of the entrant is made salient in the recruitment. Note that I do not include a version where the latter varies by caste, since making caste salient for UC entrants would not be organic and would not have

[19]Caste identity was made salient through the workers' names, but not made explicit independently. This follows the same logic as the exercise to elicit priors about out-groups in the surveys with workers and supervisors at the firm.

any parallels in commonly implemented recruitment policies. Another noteowrthy aspect of this exercise is that while the entrants are identified by name in the case of outsiders, in the referral policies only the names of the *referrers* is made salient. The implicit understanding is that respondents are able to infer the identity of the entrant from the identity of the referrer, since homophily in networks and in referral allocations is the norm in this setting.

I summarize the results from this exercise in table 13, where I report both the means of the change in expected output as a result of each type of recruitment policy, as well as pairwise comparisons between policies with fixed effects for the respondent. In column (1), I report the average change in the expected output per worker between the first session of production (which was in the information given to the respondent in figure A8), and the second session of production (for which they were asked to estimate the output). The mean for both LC outsider policies (with and without the caste label) is negative, while the other policies are positive. Respondents indicate that output will be weakly lower if the entrant is LC and is hired as an outsider, while referral-based hiring has positive impacts on output across both LC and UC hires. These comparisons are presented more robustly through pairwise tests, wherein the comparisons include fixed effects for the individual respondents.

There are three noteworthy insights from the vignettes survey. First, respondents expect LC outsiders to have a worse impact on output relative to UC outsiders (0.8 units, p-value = 0.08), when neither policy makes caste salient. In fact, LC outsider hiring is expected to lower output relative to the comparisons with all the other recruitment policies. Second, respondents expect referrals to blunt the impact of hiring more LCs. In particular, LC referrals are expected to lead to a substantially higher output per worker than LC outsiders both when caste is not made salient for either policy (1.03 units, p-value = 0.02), and when caste is made salient for both (0.75 units, p-value < 0.01). Since this specification accounts for fixed effects at the individual level, this suggests that respondents perceive referral-based hiring to substantially improve output relative to the hiring of LC outsiders, where the salient difference between the policies is simply the mode of recruitment. Third, respondents perceive caste labels to have a negative impact on output, but not statistically significantly so. An analogous comparison of the expected change in cohesion after recruitment is reported in table 14.

Lab-in-field extension

In the lab-in-field extension, workers are assigned randomly to teams and are asked to work on two sessions.

The treatment effects on output are summarized in Table 8. First, I show that LC outsider teams gain no statistically significant value in output from adding a fourth worker in the second session (column 3), and the corresponding output per worker (i.e. total units produced by the team divided by 3 in the first session and divided by 4 in the second session) decreases but not statistically significantly. The corresponding effects are large

and positive for both referral teams at the team level, while the output per worker weakly improves though not statistically significantly. Importantly, the difference-in-difference estimate comparing the treatment groups, summarized in panel B, reports statistically significant differences in output as a result of the recruitment policy. The difference-in-difference estimate in terms of total units indicates that output is 23% higher in teams where LCs are hired through referrals, relative to teams where they are hired as outsiders, holding selection fixed. The analogous effect in terms of output per worker is 21%, both statistically significant at conventional levels.

Next, I present results from the revealed behavioral measures of cooperation between workers. At the end of the second session of work, each of the four workers was paired up with a randomly chosen other worker for a symbol-matching game, where each worker was given two sheets with a random set of symbols. Their task is to find the number of symbols that are in common across both sheets. An example illustrating this task is in figure A5. Each worker plays this game over two rounds, with two randomly selected workers from their own team, drawn independently each time. The outcome for each game is a share (between 0 and 1), indicating how close the players were in absolute terms to the correct number of matches between their two sheets. In table 6, I present treatment effects on this outcome. In column 1, I report that LC referral teams show a 5 percentage point (8.3%) higher cooperation than LC outsider teams (table 5). This effect is particularly pronounced among teams working on a high coordination task, among whom cooperation is higher. Importantly, thid net difference in cooperation is evident across the team members. In columns 1 - 3 in table 6, I report effects on the games involving pairs of only incumbent workers, i.e. those who were present in the first production session. In columns 4 - 6, I report effects on the sample of games wherein one of the two players was a new worker hired for the second production session. I find similar effects on both types of teams, indicating that the mode of recruitment has a level effect on cooperation across all types of workers, not only those who are involved in the recruitment directly, either as recruits themselves or those who refer them. I show treatment effects on teams disaggregated by caste composition in table 7. In columns 1 - 3, I report treatment effects on pairs of workers who both belong to the same caste group. Here, treatment effects are generally muted, except for teams working on tasks requiring a high degree of coordination. However, I report that there are large positive effects in cross-caste pairs reported in columns 4 - 6. Pooled across all teams, cooperation is higher by 7.8 percentage points in LC referral teams relative to LC outsider teams when looking specifically at cross-caste pairs, with stronger effects of 8.8 percentage points among teams working on tasks that require a high degree of coordination.

As noted above, this design shuts down alternative explanations like selection, when comparing LC referral and LC outsiders teams. These treatment effects then suggest that the mode of recruitment has a systematic effect on morale and cohesion across the team, and particularly in terms of the quality of cross-caste links that are forged in the process of working on the production task. Similar insights are evident in these workers' survey responses.

Together, the lab-in-field extension illuminates the key mechanism driving the results

of the firm experiment. The absence of a detectable dip in output (and implicitly, cohesion) as a result of the increased diversity in treatment teams at the firm can be explained by the way in which the diversity is introduced, i.e. using the existing links between UC and LC incumbent workers to sand the edges of the diversity-element of this recruitment policy.

7 Conclusion

This paper examines the role of referral-based hiring in improving labor market access for underrepresented groups while maintaining firm productivity. Using the case of caste-based labor market stratification in India, the findings provide evidence that structured interventions in referral allocation can meaningfully expand opportunities for historically marginalized workers without imposing significant costs on firms.

The study reveals that when referral opportunities are intentionally directed toward lower-caste (LC) workers, firms see substantial gains in workforce diversity and retention. This is particularly relevant in settings where hiring is overwhelmingly driven by social networks, reinforcing occupational segregation. By reallocating referral opportunities, firms not only increase representation but also benefit from reduced turnover, which in turn fosters more stable production teams. The improved retention of LC workers appears to be a key driver of productivity gains, suggesting that referrals serve not only as a mechanism for hiring but also as a means of sustaining a committed workforce.

A central concern in the literature on diversity and firm performance is the potential trade-off between inclusion and team cohesion. Prior research suggests that diverse teams may suffer initial disruptions in productivity due to coordination challenges. However, this study finds no evidence of short-term productivity declines in treatment teams. Instead, the intervention led to higher team output, particularly in tasks requiring a high degree of coordination. These findings challenge conventional wisdom by suggesting that diversity-related frictions can be mitigated when new hires are embedded within existing networks through referrals. The additional insights from a lab-in-field extension reinforce this point: teams that integrated LC workers via referrals displayed significantly greater cooperation and social cohesion than teams where diversity was introduced through direct hiring without incumbent worker ties.

The study also sheds light on why firms and supervisors may not have previously implemented similar referral allocation strategies. The reluctance to shift away from status quo hiring practices may stem less from economic concerns about productivity and more from behavioral and institutional constraints. Supervisors tasked with allocating referrals may be guided by risk aversion and short-term performance targets, making them less willing to experiment with alternative hiring strategies even when such changes could be beneficial in the long run. In this sense, the study's findings suggest a misalignment between firm-level objectives—reducing turnover and increasing work-

force stability—and individual decision-making at the managerial level. Addressing these constraints through incentive realignment or informational interventions could encourage firms to adopt referral policies that enhance diversity while improving workforce retention.

From a policy perspective, the results of this study offer important implications for labor market interventions aimed at increasing inclusion. Many diversity-focused hiring policies rely on external mandates or quotas, which can sometimes generate resistance or unintended productivity losses. This paper provides evidence that referral-based hiring reforms offer an alternative pathway—one that leverages existing hiring mechanisms while fostering a more diverse and stable workforce. Firms operating in labor markets with strong social network effects may benefit from restructuring referral allocation in ways that broaden hiring pipelines without disrupting existing production dynamics.

Overall, these findings contribute to the broader literature on social networks in labor markets, workplace diversity, and firm efficiency. They highlight that referral-based hiring is not only a mechanism for filling vacancies but also a key determinant of workforce composition, retention, and long-term firm performance. In settings where turnover is high and recruitment frictions persist, structured referral interventions can serve as a low-cost strategy for firms to achieve both equity and efficiency goals.

Tables and figures

Table 1: Balance

Variable	(1) Control N	Mean/SE	(2) Treatment N	Mean/SE	T-test Difference (1)-(2)
Output	66	0.016	66	-0.016	0.033
		(0.125)		(0.115)	
Average team tenure	66	1.594	66	1.472	0.122
		(0.116)		(0.076)	
Share of workers with high school degrees	66	0.420	66	0.399	0.021
		(0.023)		(0.023)	
Share of workers with above median experience	66	0.494	66	0.499	-0.005
		(0.024)		(0.026)	
Self-reported team cohesion (out of 10)	66	8.811	66	8.757	0.053
		(0.071)		(0.074)	
Share of workers that matched through a referral	66	0.882	66	0.888	-0.006
		(0.019)		(0.019)	
Average number of connections at the firm before starting	66	1.809	66	1.361	0.448*
		(0.238)		(0.070)	
Share of workers that have referred anyone	66	-0.013	66	0.306	-0.319
		(0.296)		(0.028)	
F-test of joint significance (p-value)					0.598
F-test, number of observations					132

This table reports balance on a range of pre-intervention characteristics of the teams that are eventually randomized into treatment and control. Columns 1 and 2 report the number of observations and the mean and standard error at the team level for control and treatment teams respectively. Column 3 reports the magnitude of the difference between the two means, with statistical significance reported after controlling for randomization strata fixed effects, and clustering the standard errors are clustered at the team level. The p-value from the F-test of joint significance is reported in the bottom panel. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 2: Compliance with referral allocations (firm experiment)

	# LC referral hires	# LC referral hires / # referral hires
Treatment	0.240***	0.677***
	(0.058)	(0.037)
Control mean	0.288	0.241
R^2	0.12	0.51
Observations	924	467
# teams	132	126
# months	7	6

This table reports treatment effects from the firm experiment on the number of lower caste referral hires at the team × month level, and the number of lower caste referral hires as a proportion of all referral hires. The regressions use randomization strata fixed effects, and the standard errors are clustered at the team level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 3: Treatment effects on output (firm experiment)

	All teams		High coord teams		Low coord teams	
		RCT period		RCT period		RCT period
Treatment	0.108***	0.072***	0.170***	0.090**	-0.001	0.026
	(0.037)	(0.024)	(0.058)	(0.036)	(0.024)	(0.017)
Control mean	-0.016	0.004	0.043	0.084	-0.100	-0.112
High vs low p-val					0.007	0.149
R^2	0.85	0.91	0.79	0.88	0.95	0.97
Observations	5940	3432	3600	2080	2340	1352
# teams	132	132	80	80	52	52
# weeks	45	26	45	26	45	26

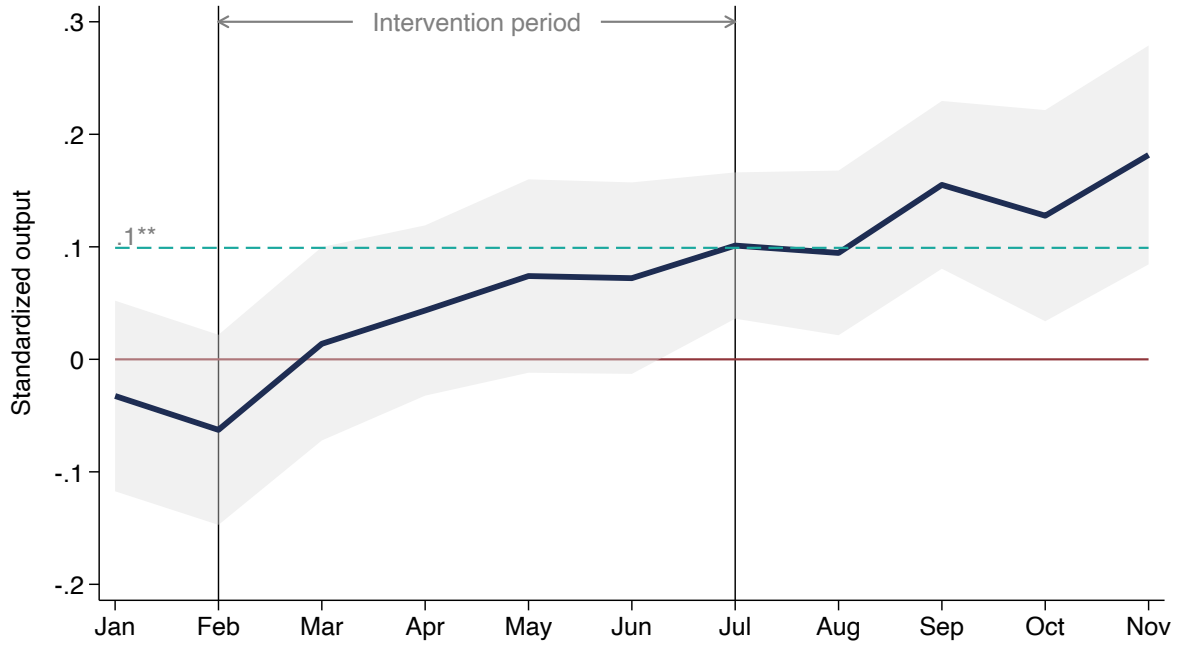
This table uses administrative production data at the team × week level. The output is standardized at the team-type level, using pre-experiment data. Columns 1 and 2 report treatment effects on the standardized output for all teams over the entire period that the data is available for (February - November, 2024) and only the period when the intervention was active, respectively. Columns 3 and 4 report treatment effects only for teams working on high coordination tasks. Columns 5 and 6 report treatment effects only for teams working on low coordination tasks. Columns 5 and 6 also report in the panel below the main point estimates p-values from testing heterogeneous treatment effects by the extent of coordination. All standard errors are clustered at the team level, and the regressions inlcude fixed effects at the randomization strata and week level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 4: Treatment effects on churn (firm experiment)

	Exits				Entries		Churn
	Total	Incumbent	New hire	LC	Total	LC	
Treatment	-0.522***	-0.074	-0.448***	0.102*	-0.626***	0.333***	-1.147***
	(0.074)	(0.057)	(0.064)	(0.054)	(0.082)	(0.047)	(0.139)
Control mean	1.472	0.684	0.788	0.331	1.281	0.190	2.753
R^2	0.26	0.49	0.21	0.07	0.24	0.14	0.16
Observations	924	924	924	924	924	924	924
# teams	132	132	132	132	132	132	132
# months	7	7	7	7	7	7	7

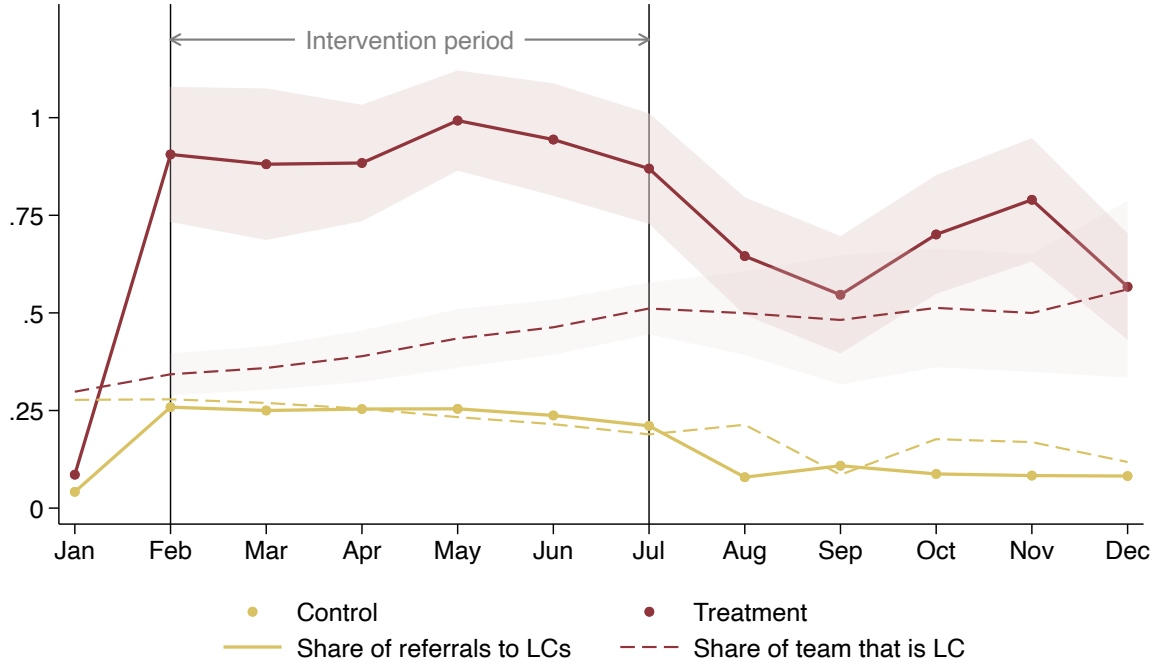
This table uses personnel data at the team × month level to report treatment effects on worker retention outcomes from the experiment at the firm. In columns 1-4, I report the treatment effect on the number of total exits, incumbent exits, exits among new hires, and exits among lower castes respectively. In columns 5-6, I report treatment effects on the number of total entries and lower caste entries respectively. Column 7 reports treatment effects on overall churn, which is the sum of total exits and total entries. All standard errors are clustered at the team level, and the regressions inlcude fixed effects at the randomization strata level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Figure 1: Dynamic treatment effects on output (firm experiment)



This figure plots the evolution of the treatment effect on a standardized measure of team output at the monthly level, with 95% confidence intervals from standard errors clustered at the team level and fixed effects at the randomization strata and week level, the latter being the unit at which the data is reported. The dashed horizontal line indicates the pooled treatment effect over the entire period, from February - November 2024.

Figure 2: Dynamic referral allocations (firm experiment)



This figure plots the evolution of the team composition and referral allocation towards lower caste members over the intervention and post-intervention period. In the pre-intervention period (January 2024), the share of lower castes was approximately 26% and the share of referrals they received was approximately 8%. In the intervention period, treatment teams allocate nearly all referrals to their lower caste members (solid maroon line), which results in a gradual increase in the share of the team that is lower caste (dashed maroon line). In this period, control teams' share of referrals allocated to lower caste workers (solid yellow line) approximately tracks the share of workers that are lower caste (dashed yellow line), since referrals are allocated randomly. After the intervention period, the share of referrals that are allocated to lower caste workers in treatment teams remains above their share of lower caste workers, while reverts to the pre-experiment rate for control teams. The shaded areas in the treatment (maroon) lines reflect 95% confidence intervals from regressions on team × month level data relative to the control (yellow) analog.

Table 5: Treatment effect on team cohesion (lab-in-field)

	All	High coord	Low coord
	(1)	(2)	(3)
LC referral teams	0.055***	0.086***	0.025
	(0.021)	(0.029)	(0.028)
UC referral teams	0.034	0.037	0.031
	(0.021)	(0.029)	(0.031)
LC outsiders mean	0.601	0.589	0.612
R^2	0.06	0.12	0.05
Observations	960	480	480

This table reports treatment effects on pairwise tests from a symbol matching game to measure the extent of team cohesion. The outcome is defined as the one minus the absolute share of true matches between the two players' sheets that are not found. Thus, the outcome ranges between 0 and 1, wherein higher values indicate more cohesion. In column 1, I report the pooled effect on all teams. Columns 2 and 3 report the point estimates for high and low coordination teams respectively. All regressions use randomization strata fixed effects, and standard errors are clustered at the team level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 6: Treatment effect on team cohesion by worker status (lab-in-field)

	Incumbent workers only			New worker involved		
	(1) All	(2) High coord	(3) Low coord	(4) All	(5) High coord	(6) Low coord
LC referral teams	0.051*	0.095**	0.007	0.060**	0.076**	0.044
	(0.028)	(0.038)	(0.037)	(0.025)	(0.036)	(0.033)
UC referral teams	0.031	0.041	0.017	0.038	0.034	0.048
	(0.030)	(0.044)	(0.039)	(0.026)	(0.032)	(0.041)
LC outsiders mean	0.610	0.588	0.632	0.592	0.590	0.594
R^2	0.08	0.19	0.07	0.13	0.17	0.18
Observations	480	240	240	480	240	240

This table reports treatment effects on pairwise tests from a symbol matching game to measure the extent of team cohesion, separately by the status of workers in the pairs. The outcome is defined as the one minus the absolute share of true matches between the two players' sheets that are not found. Thus, the outcome ranges between 0 and 1, wherein higher values indicate more cohesion. In columns 1-3, I report the effect on team cohesion considering only worker pairs where both workers were incumbent workers, i.e. were present in the first session of production prior to the recruitment. In columns 4-6, I consider only worker pairs where one of the two workers was a new worker recruited between the first and second session of production. All regressions use randomization strata fixed effects, and standard errors are clustered at the team level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 7: Treatment effect on team cohesion by worker caste (lab-in-field)

	Same-caste pairs			Cross-caste pairs		
	(1) All	(2) High coord	(3) Low coord	(4) All	(5) High coord	(6) Low coord
LC referral teams	0.028	0.079**	-0.025	0.078***	0.088**	0.062
	(0.026)	(0.035)	(0.037)	(0.028)	(0.037)	(0.040)
UC referral teams	-0.005	0.029	-0.052	0.069**	0.052*	0.096**
	(0.035)	(0.044)	(0.057)	(0.028)	(0.030)	(0.040)
LC outsiders mean	0.628	0.596	0.663	0.573	0.581	0.565
R^2	0.10	0.17	0.12	0.07	0.12	0.11
Observations	485	248	237	475	232	243

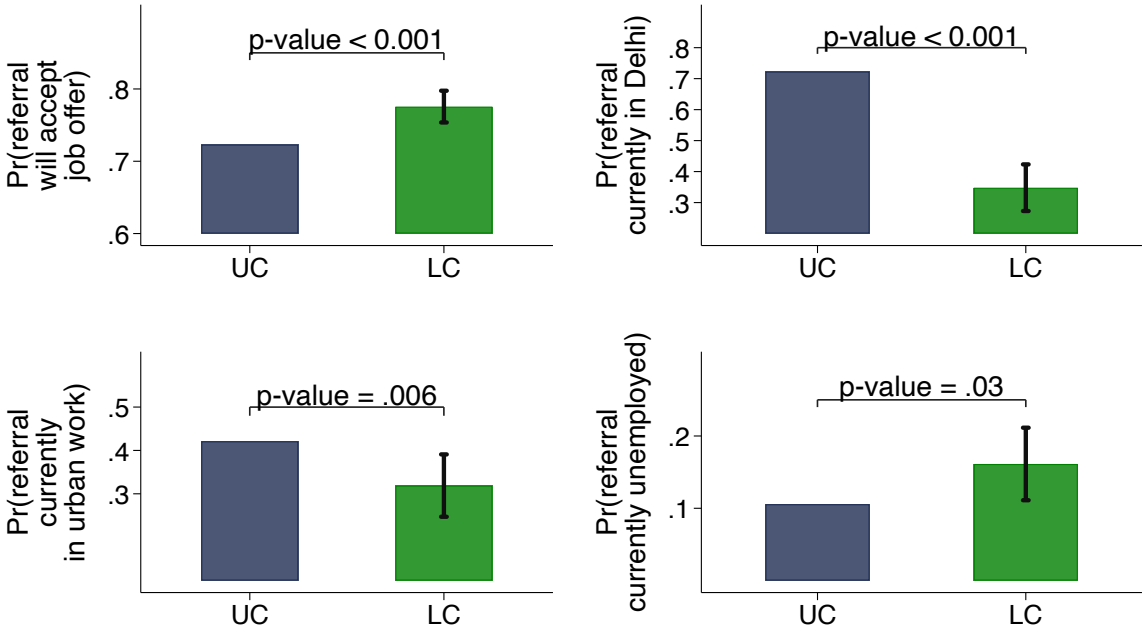
This table reports treatment effects on pairwise tests from a symbol matching game to measure the extent of team cohesion, separately by the caste group of the workers in the pairs. The outcome is defined as the one minus the absolute share of true matches between the two players' sheets that are not found. Thus, the outcome ranges between 0 and 1, wherein higher values indicate more cohesion. In columns 1-3, I report the effect on team cohesion considering only worker pairs where both workers belonged to the same caste. Note that this could include both upper caste-only pairs and lower caste-only pairs. In columns 4-6, I consider only worker pairs where both workers belonged to different caste groups. All regressions use randomization strata fixed effects, and standard errors are clustered at the team level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 8: Treatment effect on output (lab-in-field)

	All teams (N = 120)						High coordination teams (N = 75)						Low coordination teams (N = 45)					
	Total units			Units per worker			Total units			Units per worker			Total units			Units per worker		
	S1	S2	Δ	S1	S2	Δ	S1	S2	Δ	S1	S2	Δ	S1	S2	Δ	S1	S2	Δ
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
Panel A: Output by treatment group																		
β_1: LC outsider teams	54.49	63.38	9.26	18.16	15.84	-2.23	48.39	52.10	4.03	16.13	13.03	-3.01	64.50	82.17	17.67	21.50	20.54	-0.96
			(6.38)			(1.89)			(8.04)			(2.41)			(10.48)			(3.12)
β_2: LC referral teams	54.06	78.81	24.75***	18.02	19.70	1.68	48.73	74.97	26.23***	16.24	18.74	2.50	62.94	85.22	22.28***	20.98	21.31	0.32
			(4.07)			(1.29)			(5.42)			(1.68)			(6.15)			(2.01)
β_3: UC referral teams	49.88	75.00	25.13***	16.62	18.75	2.13	33.33	60.67	13.80	11.11	15.17	-0.84	47.11	91.11	44.00***	15.70	22.78	7.07
			(7.52)			(2.28)			(8.09)			(2.65)			(13.10)			(3.81)
Panel B: Difference-in-difference estimates																		
$\Delta_1 - \Delta_2$			15.49**			3.91*			22.20**			5.51*			4.61			1.28
			(7.54)			(2.28)			(9.64)			(2.65)			(13.10)			(3.81)
$\Delta_1 - \Delta_3$			15.87			4.35			9.77			2.17			26.33			8.03
			(10.43)			(3.11)			(12.63)			(3.86)			(17.41)			(5.07)
$\Delta_2 - \Delta_3$			0.37			0.44			-12.43			-3.34			21.72			6.75
			(7.82)			(2.43)			(9.56)			(3.03)			(16.61)			(4.97)

This table reports effects on output by treatment group from the lab-in-field experiment. In panel A, I report means for teams of all three recruitment types. In columns 1 and 2, I report the average units of production for each treatment group in sessions 1 and 2 respectively, with fixed effects for the workshop spell. In column 3, I report the difference in the production between the two sessions. In columns 4 and 5, I report the analog effects on units per worker, i.e. team level output divided by 3 for session 1 and divided by 4 for session 2. Column 6 then, as in column 3, reports the difference in the per-worker output between the two sessions. The same pattern of outcomes is repeated in columns 7 - 18, for high and low coordination teams respectively. In panel B, I report the difference-in-difference estimates, comparing the single differences reported in columns 3, 6, 9, 12, 15, and 18 pairwise for each of the three combinations involved across the three treatment groups. The difference-in-difference estimates use workshop spell fixed effects (as well as coordination level fixed effects in column 3), with robust standard errors. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Figure 3: Descriptive comparison of referral candidates (firm experiment)



In this figure, I report simple descriptive comparisons of the referral candidates of lower caste and upper caste incumbent workers at the firm, as reported by the referrers (i.e. the incumbent workers) themselves during the baseline survey. The first graph reports the difference in lower and upper caste referral candidates' probability of accepting the job, as reported by the referrers. The second graph reports the difference in the probability of already being in Delhi as of the baseline survey, the site of the experiment, as reported by the referrers. The third graph reports the difference in the probability of the referral candidates being in urban work as of the baseline survey, as reported by the referrers. The fourth graph reports the difference in the probability of the referral candidates being unemployed as of the baseline survey, as reported by the referrers. The p-values for these comparisons are from simple bivariate regressions by the caste of the referrer with task fixed effects and robust standard errors.

Table 9: Treatment effect on cohesion (firm experiment)

	Team cohesion (0-10)		Friends outside		Friends at work		Bargaining power	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	0.161*	-0.130	0.412	-0.037	0.002	-0.236	-0.026	-0.044
	(0.088)	(0.125)	(0.345)	(0.443)	(0.200)	(0.245)	(0.034)	(0.045)
Treatment × Lower caste		0.144		1.357*		0.973**		0.060
		(0.186)		(0.780)		(0.463)		(0.078)
Control mean	8.309	8.309	3.505	3.505	3.656	3.656	0.679	0.679
R^2	0.02	0.09	0.02	0.02	0.02	0.03	0.01	0.01
Observations	771	771	771	771	771	771	774	774

This table reports treatment effects on self-reported measures of team cohesion at the firm, as of the endline survey with workers. In columns 1-2, I report treatment effects on workers' ratings between 0 and 10 on how cohesive their team is. In columns 3-4, I report treatment effects on the number of friends they have at work with whom they would spend time outside of work. In columns 5-6, I report treatment effects on the total number of friends they have at work. In columns 7-8, I report treatment effects on whether workers believe that they have the ability to collectively bargain against the employer if a hypothetical conflict should arise. All standard errors are clustered at the team level, and the regressions inlcude fixed effects at the randomization strata level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 10: Treatment effect on job satisfaction (firm experiment)

	Value referral		Job satisfaction	
	(1)	(2)	(3)	(4)
Treatment	0.003	-0.035	0.101	0.093
	(0.035)	(0.043)	(0.099)	(0.143)
Treatment × Lower caste		0.143*		0.074
		(0.077)		(0.257)
Control mean	0.738	0.738	8.298	8.298
R^2	0.01	0.02	0.01	0.01
Observations	774	774	773	773

This table reports treatment effects on how workers value referrals and how satisfied they are at the firm. In columns 1-2, I report treatment effects on an indicator that captures whether workers value the ability to make referrals more than they would like a hypothetical 15 minute extention to their lunch break. In columns 3-4, I report treatment effects on a self-reported rating for their job satisfaction, ranging between 0 and 10. All standard errors are clustered at the team level, and the regressions inlcude fixed effects at the randomization strata level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 11: Treatment effect on exit beliefs (firm experiment)

	Expect to leave firm		Ln(Expected earnings in 2 years)		# job offers	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-0.009	-0.013	0.128	0.143	0.232	0.538
	(0.021)	(0.029)	(0.081)	(0.118)	(0.408)	(0.421)
Treatment × Lower caste		0.006		0.043		-0.603
		(0.050)		(0.190)		(0.654)
Control mean	0.888	0.888	9.863	9.863	1.043	1.043
R^2	0.01	0.01	0.06	0.06	0.05	0.05
Observations	774	774	302	302	302	302

This table reports treatment effects on self-reported measures of their expectations around exiting the firm and their outside options. In columns 1-2, I report treatment effects on whether workers expect to leave the firm in the next five years. In columns 3-4, I report treatment effects on the logged expected earnings in two years. In columns 5-6, I report treatment effects on the number of job offers the workers heard in the previous month. All standard errors are clustered at the team level, and the regressions include fixed effects at the randomization strata level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 12: Supervisors' time use

	Monitoring	Training	Recruitment	Admin work	Troubleshooting	Disciplinary	Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treatment	0.006	-0.853***	-0.159	0.356*	-0.030	0.166	0.932*
	(0.410)	(0.179)	(0.230)	(0.201)	(0.752)	(0.114)	(0.475)
Control mean	23.937	3.587	2.317	1.857	13.841	0.429	1.355
R^2	0.12	0.28	0.11	0.11	0.08	0.19	0.20
Observations	127	127	127	127	127	127	88

This table reports treatment effects on supervisors' self reported time use, as of the endline survey. Each column reports the number of hours in the previous week that supervisors spent on each of the listed tasks. All regressions inlcude fixed effects at the randomization strata level, and robust standard errors. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 13: Vignettes survey: changes in output

Recruitment type ↓→	Caste label ↓→	Change in expected output per worker after recruitment						Obs
		Mean	Pairwise test [row - column]					
			UC outsider	UC referral	LC outsider	LC referral	LC outsider	
			No	No	No	No	Yes	
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
UC outsider	No	0.16	-					93
UC referral	No	0.48	-0.09	-				107
			(0.30)					
LC outsider	No	-0.01	-0.80*	-0.52	-			112
			(0.47)	(0.35)				
LC referral	No	0.42	0.16	0.04	1.03**	-		88
			(0.27)	(0.20)	(0.43)			
LC outsider	Yes	-0.09	-0.16	-0.48*	-0.12	-0.59**	-	95
			(0.24)	(0.25)	(0.18)	(0.28)		
LC referral	Yes	0.44	0.09	0.21	0.36	-0.12	0.75***	105
			(0.25)	(0.20)	(0.24)	(0.18)	(0.26)	

This table reports average reported responses for the change in expected output per worker as a hypothetical team grows from three to four through each reported mode of recruitment. The modes are listed in the labels of rows and columns, with each cell of the matrix in columns 2-6 reporting the comparison of the columns between the mode listed in the row and the mode listed in the column. Column 1 reports the mean change in output per worker for each listed mode of recruitment as reported in the row labels. Column 7 reports the number of respondents who faced a question on the recruitment type listed in the row labels. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 14: Vignettes survey: changes in cohesion

Recruitment type ↓→	Caste label ↓→	Change in expected cohesion after recruitment						Obs
		Mean	Pairwise test [row - column]					
			UC outsider	UC referral	LC outsider	LC referral	LC outsider	
			No	No	No	No	Yes	
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
UC outsider	No	0.76	-					93
UC referral	No	0.93	0.16**	-				107
			(0.06)					
LC outsider	No	0.76	-0.02	-0.20***	-			112
			(0.05)	(0.05)				
LC referral	No	0.90	0.14**	-0.04	0.14**	-		88
			(0.06)	(0.03)	(0.05)			
LC outsider	Yes	0.61	-0.29***	-0.28***	-0.04	-0.23***	-	95
			(0.08)	(0.06)	(0.04)	(0.07)		
LC referral	Yes	0.82	0.11*	-0.13**	0.10	-0.02	0.18**	105
			(0.06)	(0.05)	(0.06)	(0.04)	(0.08)	

This table reports average reported responses for the change in expected team cohesion as a hypothetical team grows from three to four through each reported mode of recruitment. The modes are listed in the labels of rows and columns, with each cell of the matrix in columns 2-6 reporting the comparison of the columns between the mode listed in the row and the mode listed in the column. Column 1 reports the mean change in team cohesion for each listed mode of recruitment as reported in the row labels. Column 7 reports the number of respondents who faced a question on the recruitment type listed in the row labels. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

References

Afridi, Farzana, Amrita Dhillon, and Swati Sharma (May 2021). "The Ties that Bind Us: Social Networks and Productivity in the Factory". In: *CEPR Working Paper* 14687.

Afridi, Farzana, Amrita Dhillon, Shing-Yi Xin Li, et al. (May 2020). "Using Social Connections and Financial Incentives to Solve Coordination Failure: A Quasi-Field Experiment in India's Manufacturing Sector". In: *Journal of Development Economics* 144, p. 102445. DOI: 10.1016/j.jdeveco.2020.102445.

Alesina, Alberto and Eliana La Ferrara (2005). "Ethnic Diversity and Economic Performance". In: *Journal of Economic Literature* 43, pp. 762–800.

Anderson, Siwan (Jan. 2011). "Caste as an Impediment to Trade". In: *American Economic Journal: Applied Economics* 3.1, pp. 239–63. DOI: 10.1257/app.3.1.239. URL: https://www.aeaweb.org/articles?id=10.1257/app.3.1.239.

Arrow, Kenneth J. and Ron Borzekowski (2004). *Limited Network Connections and the Distribution of Wages*. Finance and Economics Discussion Series 2004-41. Board of Governors of the Federal Reserve System.

Asher, Sam, Paul Novosad, and Charlie Rafkin (Apr. 2024). "Intergenerational Mobility in India: New Measures and Estimates across Time and Social Groups". In: *American Economic Journal: Applied Economics* 16.2, pp. 66–98. DOI: 10.1257/app.20210686. URL: https://www.aeaweb.org/articles?id=10.1257/app.20210686.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul (2009). "Social Connections and Incentives in the Workplace". In: *Econometrica* 77.

Bandiera, Oriana, Imran Rasul, and Iwan Barankay (2005). "Team Incentives: Evidence from a Firm Level Experiment". In: *Journal of the European Economic Association* 3, pp. 835–845.

Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta (2013). "Cross-Country Differences in Productivity: The Role of Allocation and Selection". In: *American Economic Review* 103.1, pp. 305–334.

Beaman, Lori, Niall Keleher, and Jeremy Magruder (2018). "Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi". In: *Journal of Labor Economics*.

Beaman, Lori and Jeremy Magruder (2012). "Who gets the Job Referral? Evidence from a Social Networks Experiment". In: *American Economic Review* 102.7.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2016). "Management as a Technology?" In: *National Bureau of Economic Research Working Paper* 22327.

Bolte, Lukas, Nicole Immorlica, and Jackson (Aug. 2021). "The Role of Referrals in Immobility, Inequality, and Inefficiency in Labor Markets". In: *Working Paper*.

Brewer, Marilynn B (1999). "The psychology of prejudice: Ingroup love or outgroup hate?" In: *Journal of social issues* 55.3, pp. 429–444.

Breza, Emily and Supreet Kaur (2025). "Labor Markets in Developing Countries". In: *The Annual Review of Economics*.

Breza, Emily, Supreet Kaur, and Yogita Shamdasani (Oct. 2017). "The Morale Effects of Pay Inequality*". In: *The Quarterly Journal of Economics* 133.2, pp. 611–663. ISSN: 0033-5533. DOI: 10.1093/qje/qjx041. eprint: https://academic.oup.com/qje/article-pdf/133/2/611/24518456/qjx041.pdf. URL: https://doi.org/10.1093/qje/qjx041.

Brown, Meta, Elizabeth Setren, and Giorgio Topa (2016). "Do informal referrals lead to better matches? Evidence from a firm's employee referral system". In: *Journal of Labor Economics* 34.1, pp. 161–209.
Burks, Stephen V. et al. (2015). "The Value of Hiring through Employee Referrals". In: *The Quarterly Journal of Economics* 130.2, pp. 805–839.
Calvo-Armengol, A and Jackson (2004). "The Effects of Social Networks on Employment and Inequality". In: *American Economic Review* 94.
Caria, Stefano, Simon Franklin, and Marc Witte (2023). "Searching with Friends". In: *Journal of Labor Economics* 41.4, pp. 887–922. DOI: 10.1086/721655. eprint: https://doi.org/10.1086/721655. URL: https://doi.org/10.1086/721655.
Cassan, Guilhem, Daniel Keniston, and Tatjana Kleinberg (2022). "A Division of Laborers: Identity and Efficiency in India". In: *Working Paper*.
Chandrasekhar, Arun G., Melanie Morten, and Alessandra Peter (2020). "Network-Based Hiring: Local Benefits; Global Costs". In: *NBER Working Paper Series* 26806. DOI: 10.3386/w26806. URL: https://www.nber.org/papers/w26806.
Chetty, Raj et al. (2014). "Where is the land of opportunity? The geography of intergenerational mobility in the United States". In: *Quarterly Journal of Economics* 129.4, pp. 1553–1623.
Demirgüç-Kunt, Asli (2011). "Generating Jobs in Developing Countries: A Big Role for Small Firms". In: *World Bank Blog*. URL: https://blogs.worldbank.org/en/allaboutfinance/generating-jobs-in-developing-countries-a-big-role-for-small-firms.
Desai, Sonalde and Reeve Vanneman (2011). "India Human Development Survey: Evidence from India". In: *Economic & Political Weekly* 46.33, pp. 43–50.
Deshpande, Ashwini and Rajesh Ramachandran (Nov. 2024). "Caste disparities and affirmative action in India". In: *Oxford Review of Economic Policy* 40.3. Estimates the number of jatis in India to range from 3,000 to 6,000., pp. 630–640. DOI: 10.1093/oxrep/grae033. URL: https://academic.oup.com/oxrep/article-abstract/40/3/630/7907269.
Donovan, Kevin, Wenjie J Lu, and Todd Schoellman (2023). "Labor Market Dynamics and Development". In: *The Quarterly Journal of Economics* 138.4, pp. 2287–2325. DOI: 10.1093/qje/qjad022. URL: https://academic.oup.com/qje/article-abstract/138/4/2287/7181328.
Dustmann, Christian et al. (2016). "Referral-based Job Search Networks". In: *The Review of Economic Studies* 83.
Fershtman, Chaim and Uri Gneezy (2001). "Discrimination in a Segmented Society: An Experimental Approach". In: *Quarterly Journal of Economics* 116, pp. 351–377.
Fiske, Susan T and Shelley L Neuberg (1990). "A continuum model of impression formation: From category-based to individuating processes". In: *Advances in experimental social psychology*. Vol. 23. Academic Press, pp. 1–74.
Ghosh, Arkadev (Dec. 2023). "Religious Divisions and Production Technology: Experimental Evidence from India". In: *Working Paper*.
Granovetter, Mark S (1973). "The strength of weak ties". In: *American journal of sociology* 78.6, pp. 1360–1380.
Group, Manpower (2024). *Talent Shortage Survey*. https://go.manpowergroup.com/talent-shortage.

Hall, Robert E. and Charles I. Jones (1999). "Why Do Some Countries Produce So Much More Output per Worker than Others?" In: *The Quarterly Journal of Economics* 114.1, pp. 83–116.

Halvorsen, Sandra Kristine (2021). *Labour Turnover and Workers' Well-being in the Manufacturing Industry in Ethiopia*. Working Paper 2021/36. UNU-WIDER Working Paper Series. UNU-WIDER. URL: https://www.wider.unu.edu/sites/default/files/Publications/Working-paper/PDF/wp2021-36-labour-turnover-workers-well-being-manufacturing-industry-Ethiopia.pdf.

Hamilton, Barton H., Jack A. Nickerson, and Hideo Owan (2003a). "Team Incentives and Worker Heterogeneity: An Empirical Analysis of Tournaments and Productivity". In: *Journal of Political Economy* 111.3, pp. 465–497.

– (2003b). "Team Incentives and Worker Heterogeneity: An Empirical Analysis of Tournaments and Productivity". In: *Journal of Political Economy* 111.3, pp. 465–497.

Heath, Rachel (2018). "Why do Firms Hire Through Referrals?" In: *Journal of Political Economy* 126.4.

Hjort, Jonas (Oct. 2014). " Ethnic Divisions and Production in Firms *". In: *The Quarterly Journal of Economics* 129.4, pp. 1899–1946. ISSN: 0033-5533. DOI: 10.1093/qje/qju028. eprint: https://academic.oup.com/qje/article-pdf/129/4/1899/30699142/qju028.pdf. URL: https://doi.org/10.1093/qje/qju028.

Hnatkovska, Viktoria, Amartya Lahiri, and Sourabh B. Paul (2013). "Breaking the Caste Barrier: Intergenerational Mobility in India". In: *Journal of Human Resources* 48.

Huckman, Robert S. and Bradley R. Staats (2013). "Fluid Tasks and Fluid Teams: The Impact of Diversity in Experience and Team Familiarity on Team Performance". In: *Manufacturing & Service Operations Management* 15.1, pp. 72–91.

Ichniowski, Casey, Kathryn Shaw, and Giovanna Prennushi (1997). "The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines". In: *American Economic Review* 87.3, pp. 291–313.

ILO (2018). *World Employment and Social Outlook: Trends 2018*. Geneva: International Labour Organization.

– (2024). *World Employment and Social Outlook: Trends 2024*. Geneva: International Labour Organization. URL: https://www.ilo.org/wcmsp5/groups/public/---dgreports/---dcomm/---publ/documents/publication/wcms_834081.pdf.

Jovanovic, Boyan (1979). "Job Matching and the Theory of Turnover". In: *Journal of Political Economy* 87.5, Part 1, pp. 972–990. DOI: 10.1086/260808.

Kain, John F. (1968). "Housing Segregation, Negro Employment, and Metropolitan Growth". In: *Quarterly Journal of Economics* 82.2, pp. 371–394.

Kugler, Adriana (2003). "Employee Referrals and Efficiency Wages". In: *Labor Economics* 10.

Kumar, Naveen and Kartik Srivastava (2025). "The long-run effects of targeting schooling investments at historically disadvantaged groups". In: *Working Paper*.

Lazear, Edward P. (2000). "Performance Pay and Productivity". In: *American Economic Review* 90, pp. 1346–1361.

Lowe, Matt (June 2021). "Types of Contact: A Field Experiment on Collaborative and Adversarial Caste Integration". In: *American Economic Review* 111.6. Estimates the effects of collaborative and adversarial intergroup contact among Indian men from different castes through a cricket league experiment., pp. 1807–1844. DOI: 10.1257/

aer.20191780. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20191780.
Mas, Alexandre and Enrico Moretti (2009). "Peers at Work". In: *American Economic Review* 99, pp. 112–145.
Montgomery, D (1991). "Social Networks and Labor-Market Outcomes: Toward an Economic Analysis". In: *American Economic Review* 81.
Moorjani, Priya et al. (Sept. 2013). "Genetic Evidence for Recent Population Mixture in India". In: *The American Journal of Human Genetics* 93.3. This study provides genetic evidence indicating that endogamy became prevalent in Indian society approximately 2,000 years ago, reinforcing the caste system., pp. 422–438. DOI: 10.1016/j.ajhg.2013.07.006. URL: https://www.cell.com/ajhg/fulltext/S0002-9297(13)00383-X.
Moretti, Enrico (2012). *The New Geography of Jobs*. Boston: Houghton Mifflin Harcourt.
Munshi, Kaivan (2003). "Networks in the Modern Economy: Mexican Migrants in the U.S. Labor Market". In: *The Quarterly Journal of Economics* 118.2, pp. 549–599.
– (Dec. 2019). "Caste and the Indian Economy". In: *Journal of Economic Literature* 57.4. This article examines the pervasive role of caste in various aspects of India's economic life, including education, labor markets, and access to public resources., pp. 781–834. DOI: 10.1257/jel.20171307. URL: https://www.aeaweb.org/articles?id=10.1257/jel.20171307.
Munshi, Kaivan and Mark Rosenzweig (2015). "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap". In: *American Economic Review* 105.1, pp. 46–50.
Oh, Suanna (Aug. 2023). "Does Identity Affect Labor Supply?" In: *American Economic Review* 113.8. Investigates how caste identity influences job-specific labor supply in rural India through a field experiment., pp. 2055–2083. DOI: 10.1257/aer.20211826. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20211826.
Osman, Adam et al. (2021). "Connections, Referrals, and Hiring Outcomes: Evidence from an Egyptian Retail Sector Survey". In: *Working Paper*. URL: http://www.adam-osman.com.
Ottaviano, Gianmarco IP and Giovanni Peri (2006). "The economic value of cultural diversity: Evidence from US cities". In: *Journal of Economic Geography* 6.1, pp. 9–44.
Pallais, Amanda and E Sands (2016). "Why the Referential Treatment? Evidence from Field Experiments on Referrals". In: *Journal of Political Economy* 124.6.
Patacchini, Eleonora and Yves Zenou (2012). "Ethnic Networks and Employment Outcomes". In: *Regional Science and Urban Economics* 42.6, pp. 938–949.
Prat, Andrea (2002). "Should a Team Be Homogeneous?" In: *European Economic Review* 46, pp. 1187–1207.
Roethlisberger, Fritz J. and William J. Dickson (1939). *Management and the Worker*. Cambridge: Harvard University Press.
Sapin, Marlène et al. (2020). "The ISSP 2017 social networks and social resources module". In: *International Journal of Sociology* 50.1, pp. 1–25. DOI: 10.1080/00207659.2020.1712157.
Schmutte, Ian (2015). "Job Referral Networks and the Determination of Earnings in Local Labor Markets". In: *Journal of Labor Economics* 33.1, pp. 1–32. DOI: 10.1086/677389.

Schofield, Heather (2014). *The Economic Costs of Low Caloric Intake: Evidence from India*. Working Paper.

Shiferaw, Admasu, Måns Söderbom, and Getnet Alemu Zewdu (2020). *Job Flows and Worker Turnover: Evidence from Matched Employer-Employee Data from Ethiopia*. Tech. rep. Working Paper F-32420-ETH-11. International Growth Centre. URL: https://www.theigc.org/sites/default/files/2021/04/Shiferaw-et-al-April-2020-Working-paper.pdf.

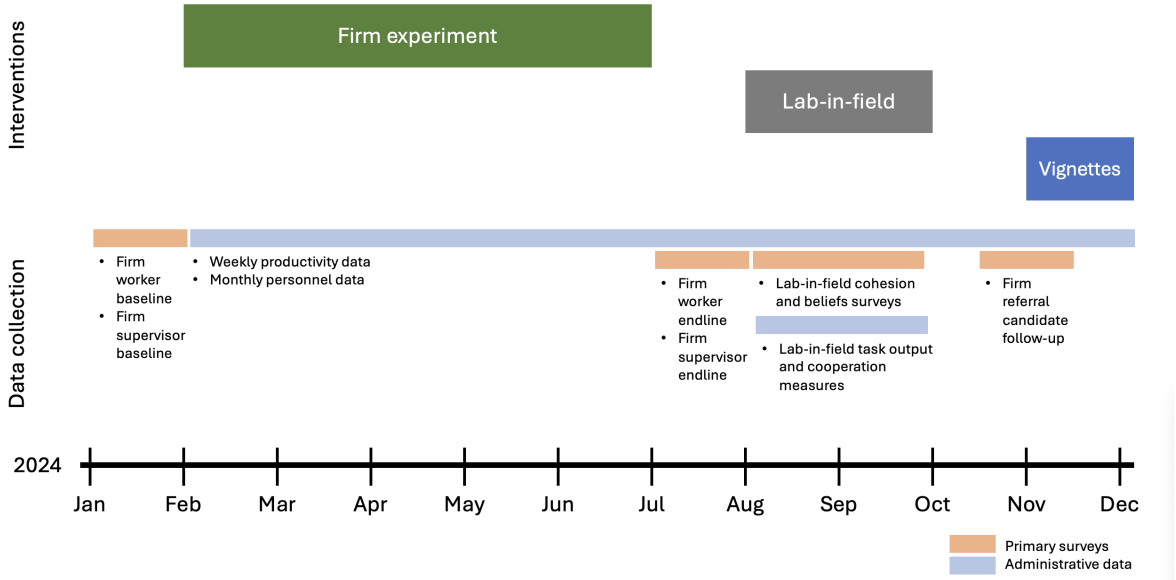
Skelton, Angie R., Deborah Nattress, and Rocky J. Dwyer (2019). "Predicting manufacturing employee turnover intentions". In: *Journal of Economics, Finance and Administrative Science* 25.49. Open Access, Article publication date: 4 June 2019, Issue publication date: 15 June 2020. ISSN: 2218-0648. DOI: 10.1108/jefas-07-2018-0069.

Tybout, James (2000). "Manufacturing Firms in Developing Countries: How Well Do They Do, and Why?" In: *Journal of Economic Literature* 38.1, pp. 11–44.

Witte, Marc (2021). "Why do workers make job referrals? Experimental Evidence from Ethiopia". In: *Working Paper*.

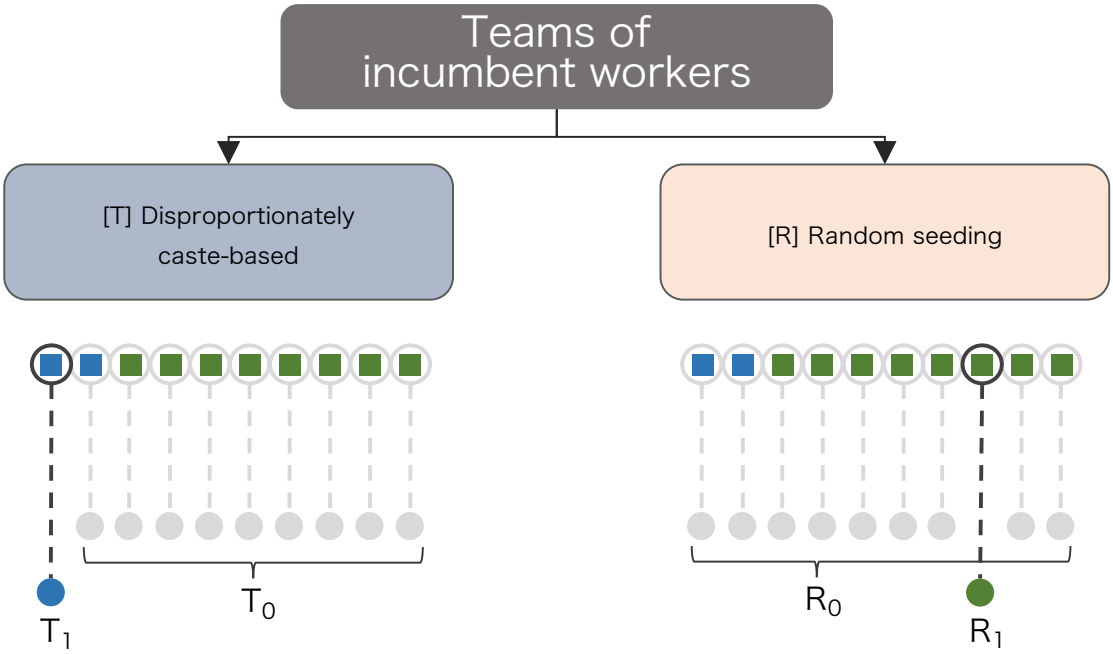
Appendix A: Study implementation

Figure A1: Study timeline



This figure plots the timeline for the study, noting both the different experiments as well as the rounds of data collection over the year of 2024.

Figure A2: Fim experiment design



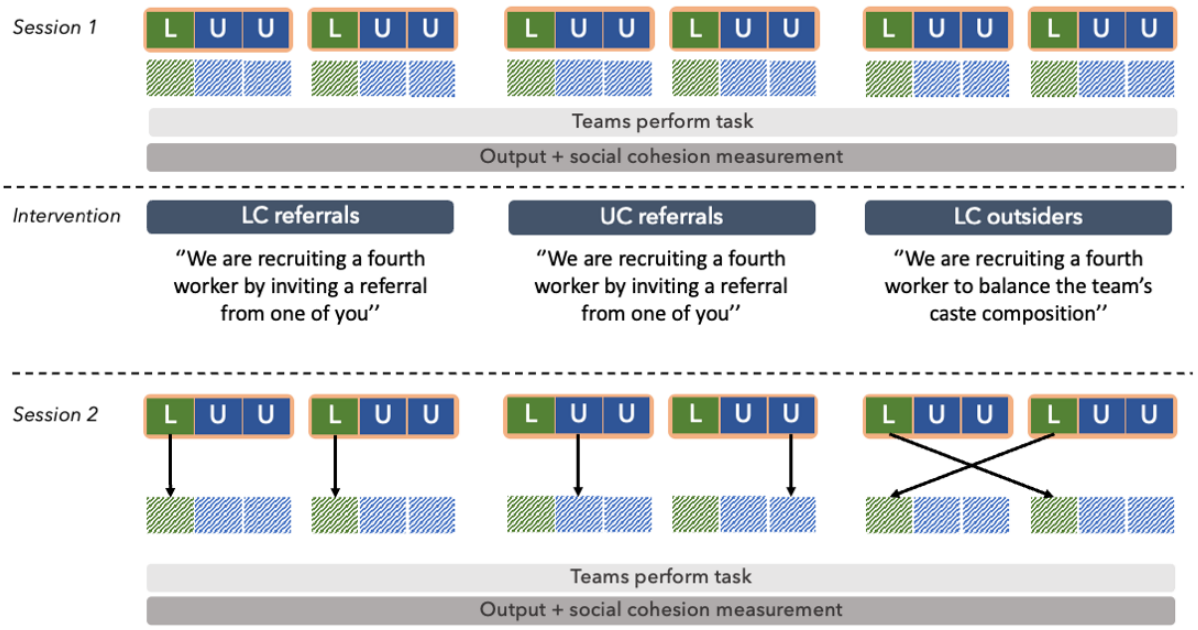
This figure plots the design of the firm intervention. The square icons represent incumbent workers, with green squares representing UCs and blue squares representing LCs. At the baseline at the firm, each incumbent worker lists (an unconstrained number of) referral candidates, represented here by circles. In the treatment group, an LC incumbent worker (blue square) gets the opportunity to refer someone. This worker's referral candidate is T_1 (blue circle, to indicate homophily such that referrals and referrers belong to the same group, though that is not mandated in the referral elicitation). T_0 represents the candidates of the other incumbent workers who are not called upon to refer from their network. In the control group, the referral opportunity goes to a UC incumbent worker (green square), chosen at random. This worker's referral candidate is R_1 (green circle). R_0 represents the candidates of the other incumbent workers.

Figure A3: Referral allocation lists

Line number	Department	Shift number	Order of asking	Employee ID
1	ASSEMBLY	A	1	[redacted]12
1	ASSEMBLY	B	1	[redacted]58
1	ASSEMBLY	B	2	[redacted]00
1	ASSEMBLY	C	1	[redacted]25
1	CUTTING	A	1	[redacted]14
1	CUTTING	A	2	[redacted]85
1	CUTTING	A	3	[redacted]24
1	CUTTING	B	1	[redacted]07
1	CUTTING	B	2	[redacted]22
1	CUTTING	C	1	[redacted]54
1	CUTTING	C	2	[redacted]55
1	MAINTENANCE	A	1	[redacted]74
1	MAINTENANCE	A	2	[redacted]75
1	MAINTENANCE	C	1	[redacted]46
1	MAINTENANCE	C	2	[redacted]19
1	PACKING	A	1	[redacted]49
1	PACKING	A	2	[redacted]10
1	PACKING	B	1	[redacted]41
1	PACKING	B	2	[redacted]09
1	PACKING	B	3	[redacted]20

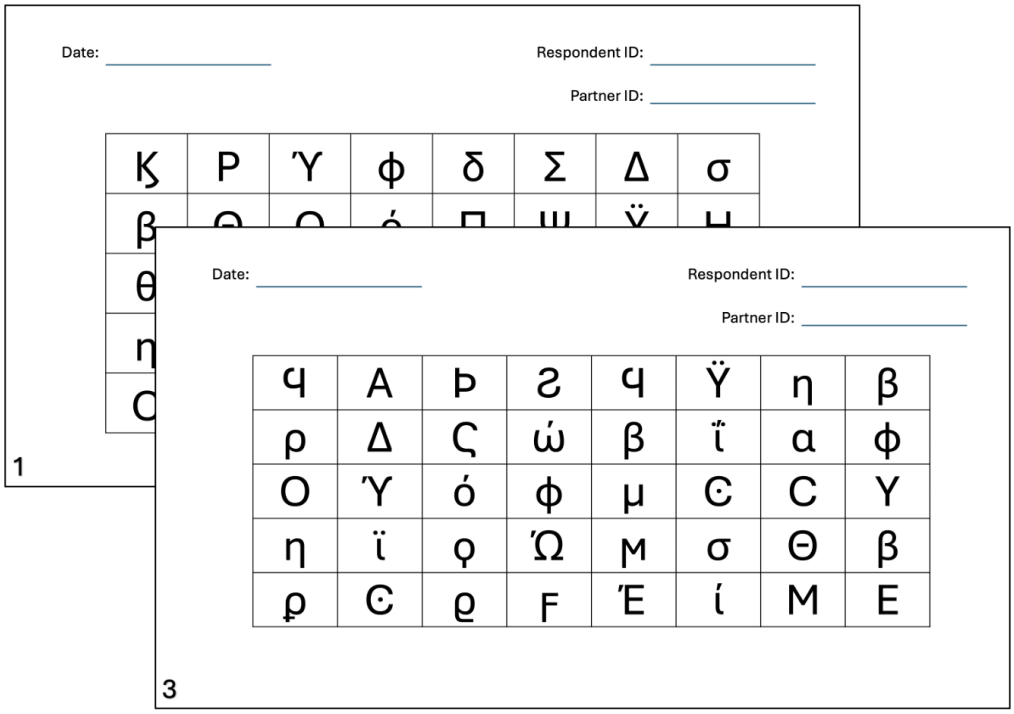
This figure shows an example of the spreadsheets shared with the firm's HR department each month, listing the referral allocations they are asked to make in order. Note that the order is listed for both treatment and control teams alike, since both sets of supervisors are uniformly asked to defer to this list.

Figure A4: Lab-in-field experiment design



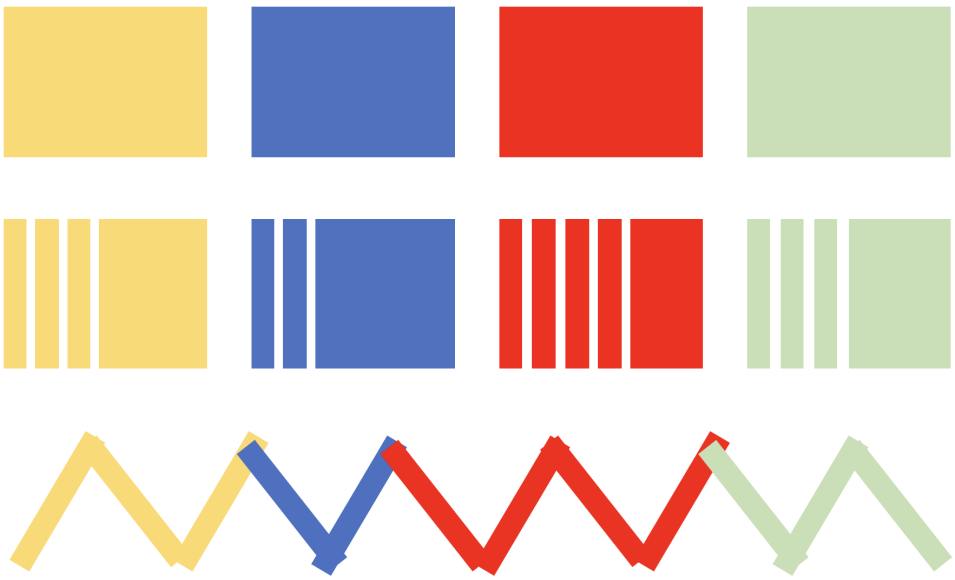
This figure shows the design of the lab-in-field experiment. Each of the workers recruited for this workshop are asked to identify referral candidates. In the first session of the lab-in-field exercise, randomly formed teams with one LC and two UCs each are asked to perform a task. After the production, each team is sorted into one of three recruitment policies, with the message given to incumbents detailed in the figure. In this recruitment process, each team increases from three to four workers. In the second session, these teams of four workers each work on another iteration of the same task. After both production sessions, the output is measured and all the workers are surveyed.

Figure A5: Symbol matching game



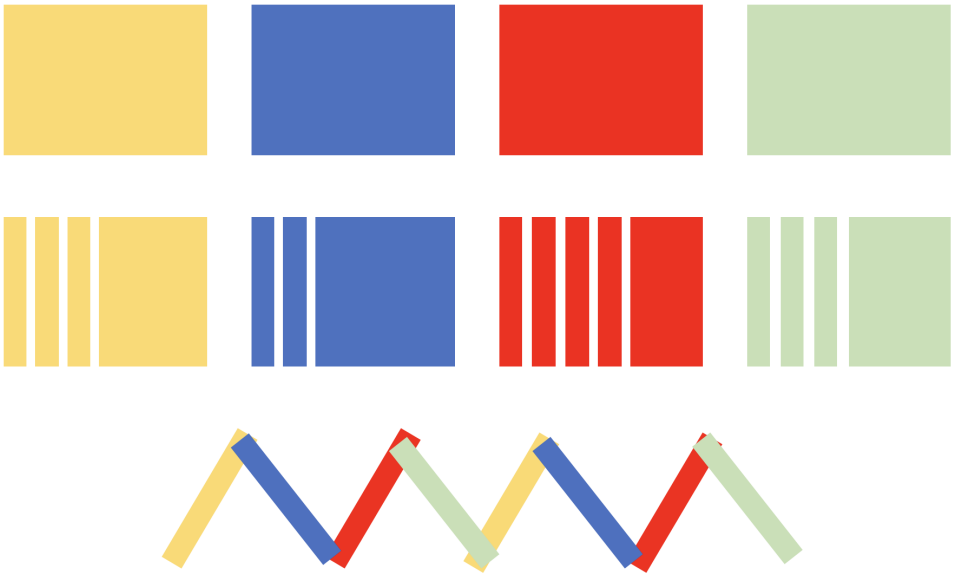
This figure shows an example of the sheets used by the workers in the symbol matching exercise used to get revealed measures of social cohesion.

Figure A6: Illustrative example of the lab-in-field task: low coordination teams



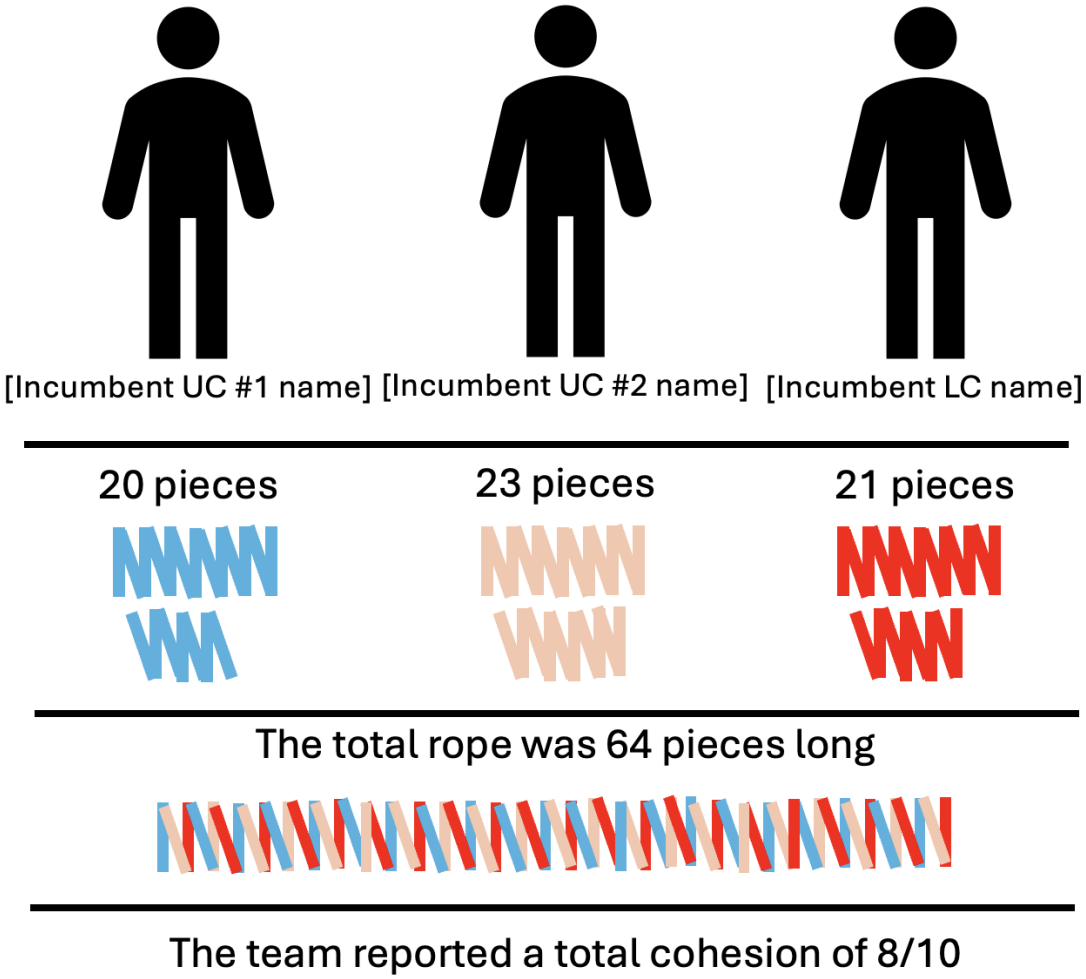
This figure represents the task that the workers in the lab-in-field exercise worked on. Each worker was given a sheet of fabric of a predetermined size and color, represented in the first row of this figure. They were then asked to cut strips of predetermined dimensions, represented in the second row. Finally, the team was asked to tie together these strips, represented in the third row. The team was compensated for the length of this combined rope, made up of the strips tied toegether. In low-coordination teams, the task did not require for any particular order of colors to be followed in constructing the rope. The incentive and effort compatible thing for workers in these teams to do is to cut and tie as many strips as possible individually, and simply to tie their individual ropes together to make a cumulative rope that is as long as possible.

Figure A7: Illustrative example of the lab-in-field task: high coordination teams



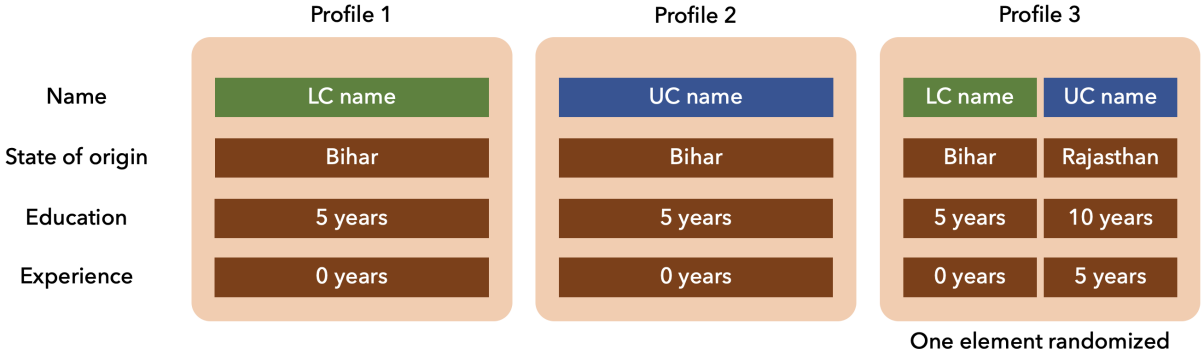
This figure represents the task in the lab-in-field similar to the previous figure, but for high-coordination teams. The key difference between these team types (and thus figures) is that high-coordination teams were asked to follow a pre-determined order of colors in constructing the rope from the strips. This necessitated workers to coordinate and time their effort together in order to construct a rope that is as long as possible.

Figure A8: Information provided in the vignettes survey



This figure was used in the vignette survey to make concrete for respondents the team for which a range of recruitment policies was then presented.

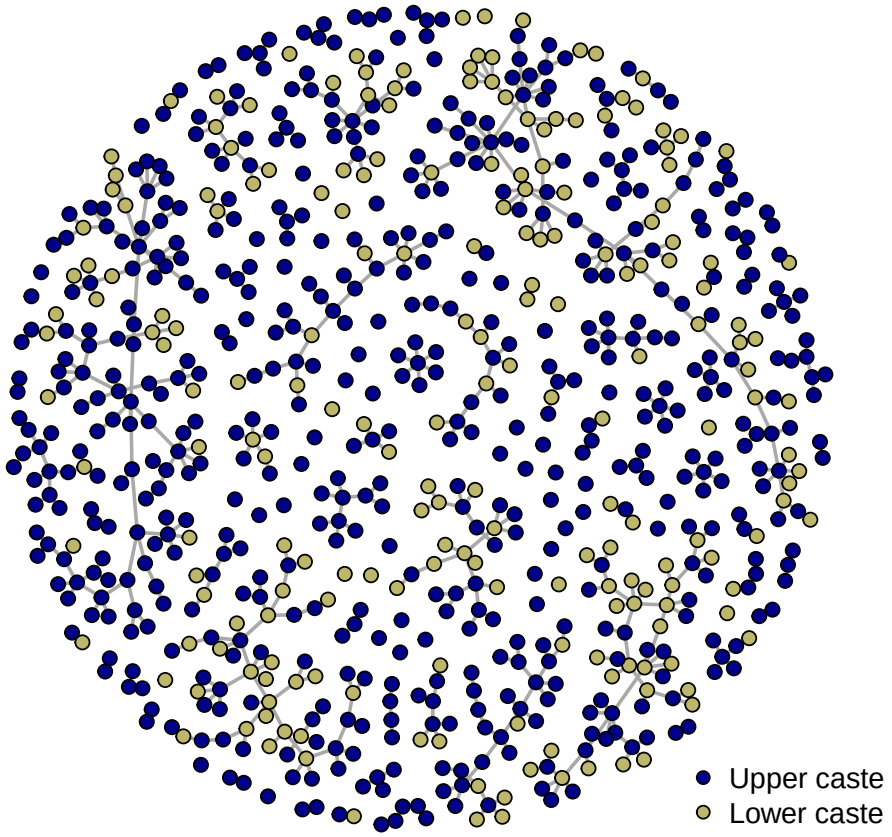
Figure A9: Hypothetical profiles for belief elicitation



This figure represents the three profiles of hypothetical workers used in the baseline and endline surveys at the firm to elicit worker and supervisor beliefs about out-groups.

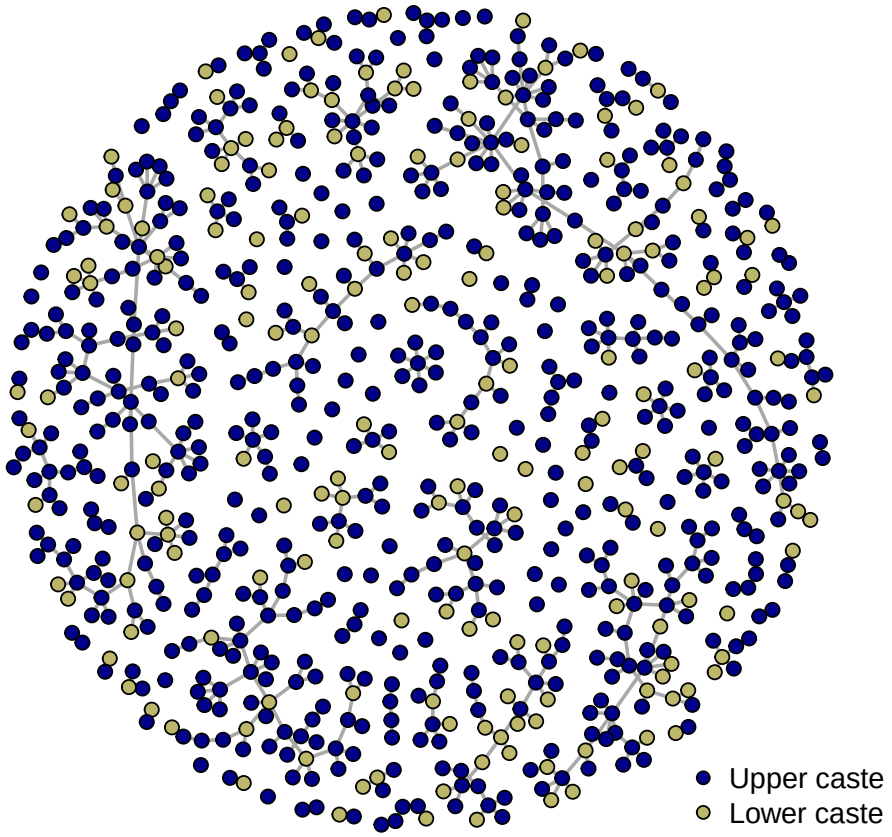
Appendix B: Supplementary exhibits

Figure B1: Incumbent worker network at baseline



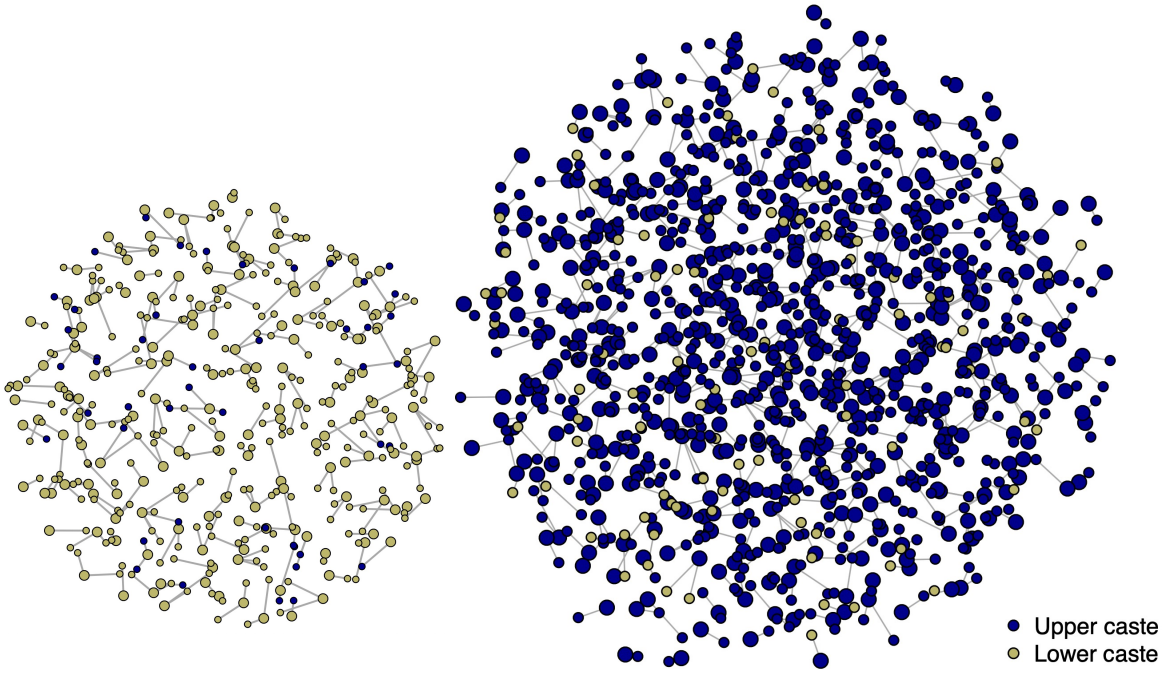
This figure represents a network graph by caste for incumbent workers at the firm at the time of the baseline, with links between referrers and referrals.

Figure B2: Incumbent worker network at baseline (if random)



This figure represents a benchmark network graph relative to the previous slide, representing what links would look like if referrals were made proportionately by caste based on their share in the workforce.

Figure B3: Network of referral elicitation



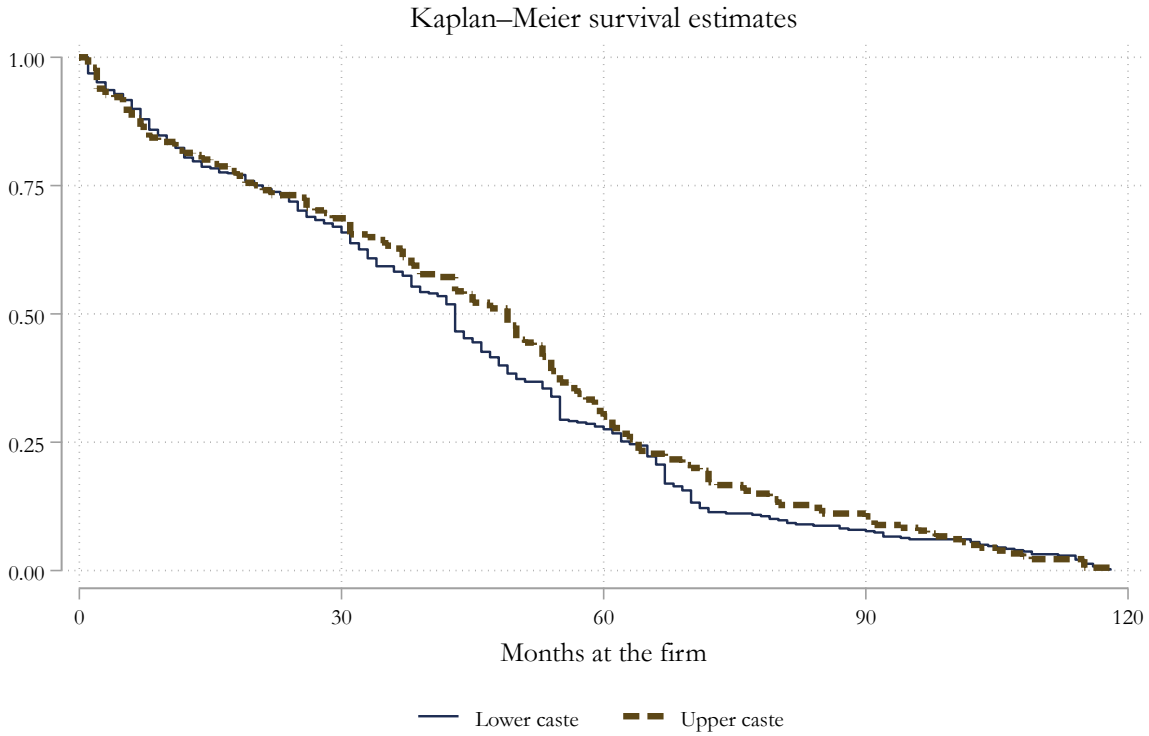
This figure represents a network graph focusing on the links between incumbent workers and their referral candidates, elicited at baseline. The graphs are represented separately for UC and LC incumbent workers, with each incumbent worker listing an average of 2.2 referral candidates.

Table B1: Firm endline: team cohesion HTEs

	Team cohesion (0-10)		Friends outside		Friends at work		Bargaining power	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	-0.130	-0.052	-0.037	-0.178	-0.236	-0.171	-0.044	-0.045
	(0.125)	(0.145)	(0.443)	(0.536)	(0.245)	(0.296)	(0.045)	(0.056)
Lower caste	0.722***	0.501***	-0.799*	-0.976*	-0.850***	-0.580	-0.044	0.073
	(0.128)	(0.191)	(0.459)	(0.539)	(0.313)	(0.485)	(0.064)	(0.104)
More diverse at baseline		0.123		-0.020		0.102		0.044
		(0.187)		(0.569)		(0.453)		(0.062)
Treatment × Lower caste	0.144	0.273	1.357*	2.123*	0.973**	1.279*	0.060	-0.028
	(0.186)	(0.238)	(0.780)	(1.099)	(0.463)	(0.694)	(0.078)	(0.118)
Treatment × More diverse at baseline		-0.205		0.344		-0.150		0.003
		(0.292)		(0.941)		(0.585)		(0.099)
Lower caste × Lower caste		0.345		0.288		-0.515		-0.205
		(0.262)		(0.854)		(0.633)		(0.128)
Triple interaction		-0.146		-1.405		-0.486		0.145
		(0.398)		(1.589)		(0.953)		(0.157)
Control mean	8.309	8.309	3.505	3.505	3.656	3.656	0.679	0.679
R^2	0.09	0.10	0.02	0.02	0.03	0.04	0.01	0.01
Observations	771	771	771	771	771	771	774	774

This table reports heterogeneous treatment effects on self reported cohesion, with respect to initial team diversity and the caste group of the respondents. All standard errors are clustered at the team level, and the regressions inlcude fixed effects at the randomization strata level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Figure B4: Baseline retention rates by caste



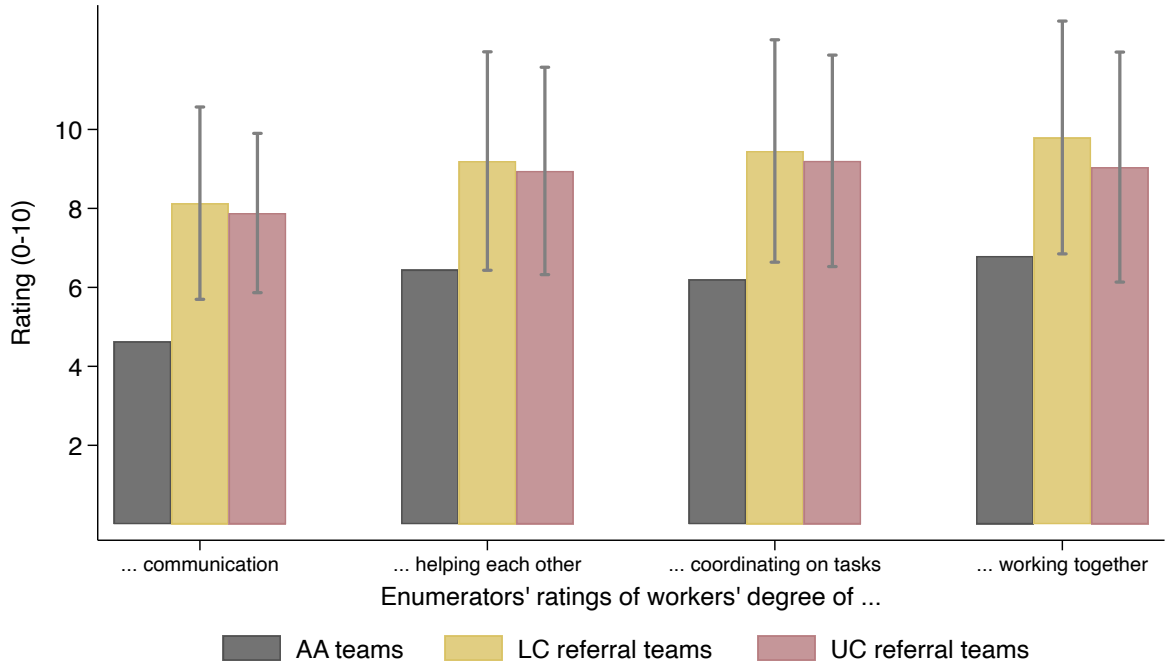
This figure shows estimates from a survival analysis among the set of incumbent workers at the firm as of the baseline, showing that lower caste workers have a higher expected retention rate.

Table B2: Lab-in-field experiment: perceived fairness

	Fairness of recruitment (0-10)			Satisfaction with workshop (0-10)		
	(1) All	(2) High coord	(3) Low coord	(4) All	(5) High coord	(6) Low coord
LC referral teams	-0.552***	-0.436**	-0.731***	-0.267**	-0.311**	-0.191
	(0.149)	(0.191)	(0.187)	(0.129)	(0.141)	(0.191)
UC referral teams	-0.791***	-0.921***	-0.565**	-0.061	-0.603***	0.864***
	(0.240)	(0.336)	(0.279)	(0.174)	(0.207)	(0.191)
LC outsiders mean	9.075	9.112	9.014	9.285	9.500	8.929
R^2	0.10	0.14	0.13	0.06	0.09	0.20
Observations	471	295	176	471	295	176

This table reports treatment effects on the perceived fairness of recruitment and satisfaction with the workshop, as reported in the endline survey at the lab-in-field experiment. Standard errors are clustered at the team level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Figure B5: Enumerators' observations of worker cohesion



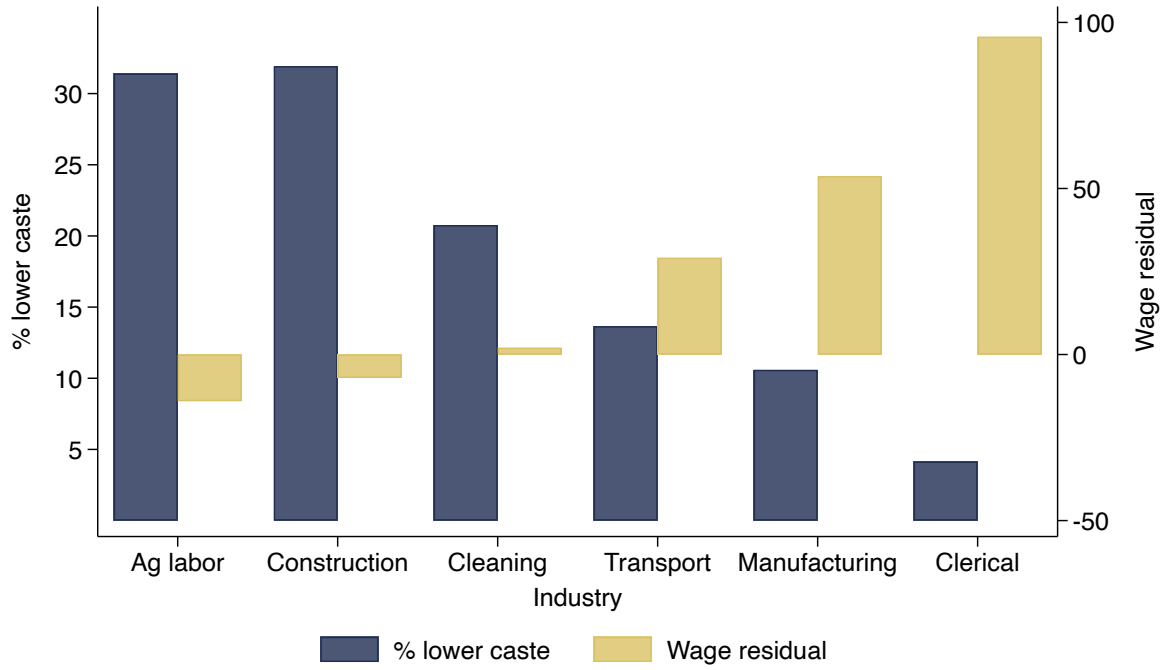
This figure represents a bar graph plotting the extent of cohesion across the three types of teams in the lab-in-field exercise, as observed by the enumerators and rated on a scale of 0 - 10.

Table B3: Lab-in-field experment: beliefs about out-groups

	Good at job			Hardworking			Good team member		
	(1) All	(2) High coord	(3) Low coord	(4) All	(5) High coord	(6) Low coord	(7) All	(8) High coord	(9) Low coord
LC referral teams	0.055	0.121	-0.097	-0.277	-0.211	-0.364	0.291*	0.425*	0.040
	(0.152)	(0.178)	(0.305)	(0.180)	(0.219)	(0.304)	(0.174)	(0.216)	(0.216)
UC referral teams	-0.080	-0.001	-0.242	-0.269*	-0.215	-0.282	0.093	0.090	0.113
	(0.162)	(0.205)	(0.274)	(0.151)	(0.215)	(0.226)	(0.169)	(0.211)	(0.190)
LC outsiders mean	0.125	0.070	0.220	0.140	0.081	0.240	-0.294	-0.372	-0.160
R^2	0.03	0.04	0.04	0.05	0.05	0.08	0.07	0.14	0.10
Observations	349	220	129	349	220	129	349	220	129

This table reports treatment effects on beliefs about hypothetical lower caste workers, as reported in the endline survey at the lab-in-field experiment. Standard errors are clustered at the team level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Figure B6: Descriptive comparison of worker caste composition and industry premia



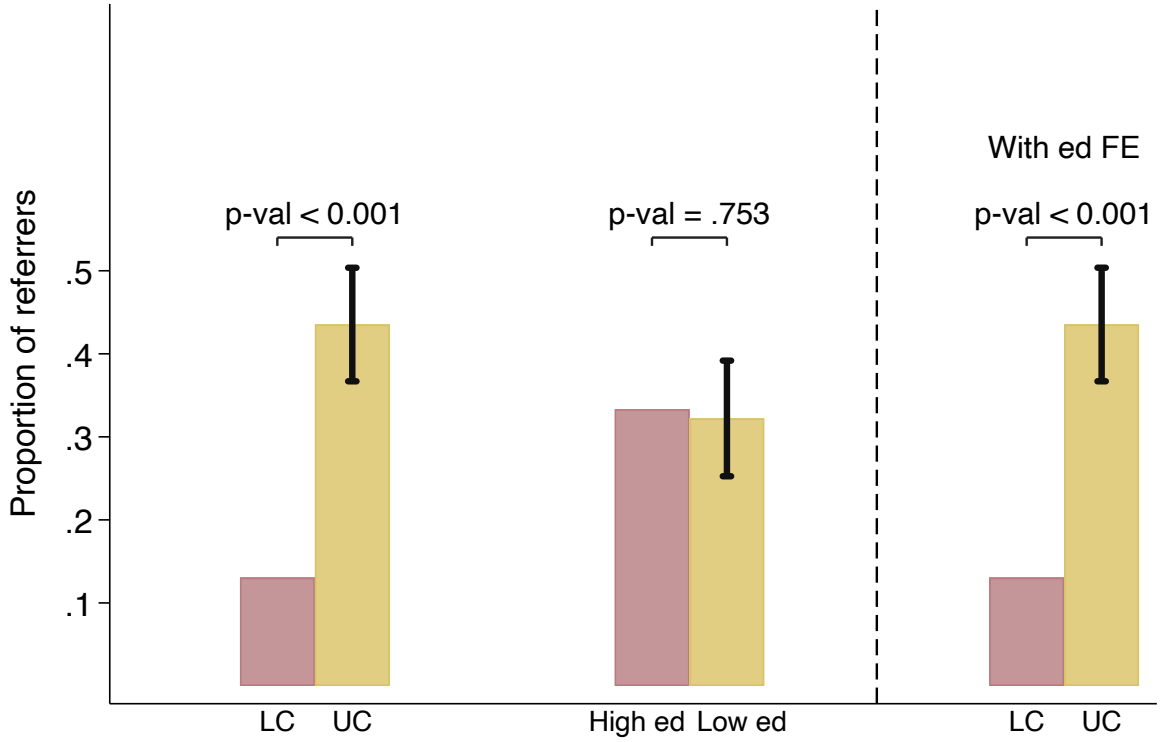
This figure represents original analysis of the nationally representative Indian Human Development Survey 2014-15 data. On the right axis, I plot the wage residual across the lowest key industry-occupation groups reported in the survey, controlling for age, education, gender, and district of birth. On the left axis, I plot the share of the workforce in the given industry-occupation group that is lower caste.

Table B4: Lab-in-field experment: Willingness to connect with teammates

	Advice			Socialize			Borrowing		
	(1) All	(2) H	(3) L	(4) All	(5) H	(6) L	(7) All	(8) H	(9) L
LC referral teams	-0.000***	-0.000	0.000	-0.000	-0.000	0.000	0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
UC referral teams	0.000	-0.000***	0.000	0.000*	0.000**	0.000	0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
LC outsiders mean	0.647	0.674	0.604	0.333	0.388	0.245	0.432	0.407	0.472
R^2	0.99	1.00	0.97	1.00	1.00	1.00	0.99	1.00	0.97
Observations	352	221	131	352	220	132	354	221	133

This table reports treatment effects on workers' future willingness to meet and socialize with their teammates, as reported in the endline survey at the lab-in-field experiment, along the three dimensions of social interactions listed: seeking advice, socializing, and borrowing money. Standard errors are clustered at the team level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Figure B7: Targeting of referral opportunities at baseline



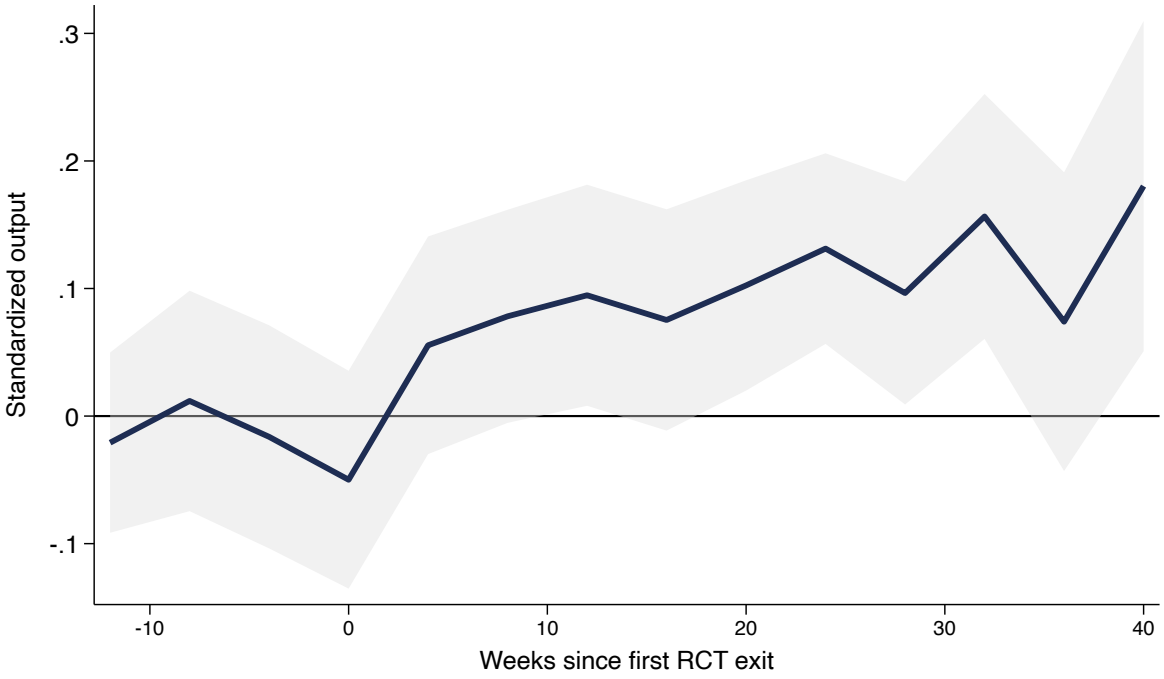
This figure uses data from the baseline survey at the firm to compare the referral allocation across UC and LC workers at the firm by the supervisors, with and without education FEs, which is used as a proxy for individual ability absent any other data.

Table B5: Lab-in-field experiment: Hypothetical referral allocation after lab-in-field

	Overall			Fairness			Productivity			Retention		
	(1) LC min	(2) Equal	(3) LC maj	(4) LC min	(5) Equal	(6) LC maj	(7) LC min	(8) Equal	(9) LC maj	(10) LC min	(11) Equal	(12) LC maj
LC referral teams	-0.007	-0.099	-0.153**	0.027	-0.048	-0.128**	-0.018	0.029	-0.105*	-0.013	-0.010	-0.100*
	(0.056)	(0.065)	(0.061)	(0.059)	(0.056)	(0.058)	(0.070)	(0.069)	(0.059)	(0.065)	(0.057)	(0.059)
UC referral teams	-0.241***	-0.034	0.034	-0.255***	-0.001	0.075	-0.127*	0.081	0.094	-0.069	0.065	-0.039
	(0.068)	(0.069)	(0.079)	(0.070)	(0.065)	(0.075)	(0.076)	(0.076)	(0.070)	(0.081)	(0.077)	(0.079)
LC outsiders mean	0.677	0.711	0.675	0.592	0.711	0.702	0.492	0.586	0.719	0.492	0.656	0.754
R^2	0.07	0.04	0.08	0.09	0.02	0.06	0.07	0.03	0.04	0.03	0.02	0.03
Observations	318	309	315	318	309	315	318	309	315	318	309	315

This table reports measures of how workers at the lab-in-field experiment would allocate future referral opportunities among a hypothetical set of workers. Standard errors are clustered at the team level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Figure B8: Treatment effect on output with respect to retention (firm experiment)



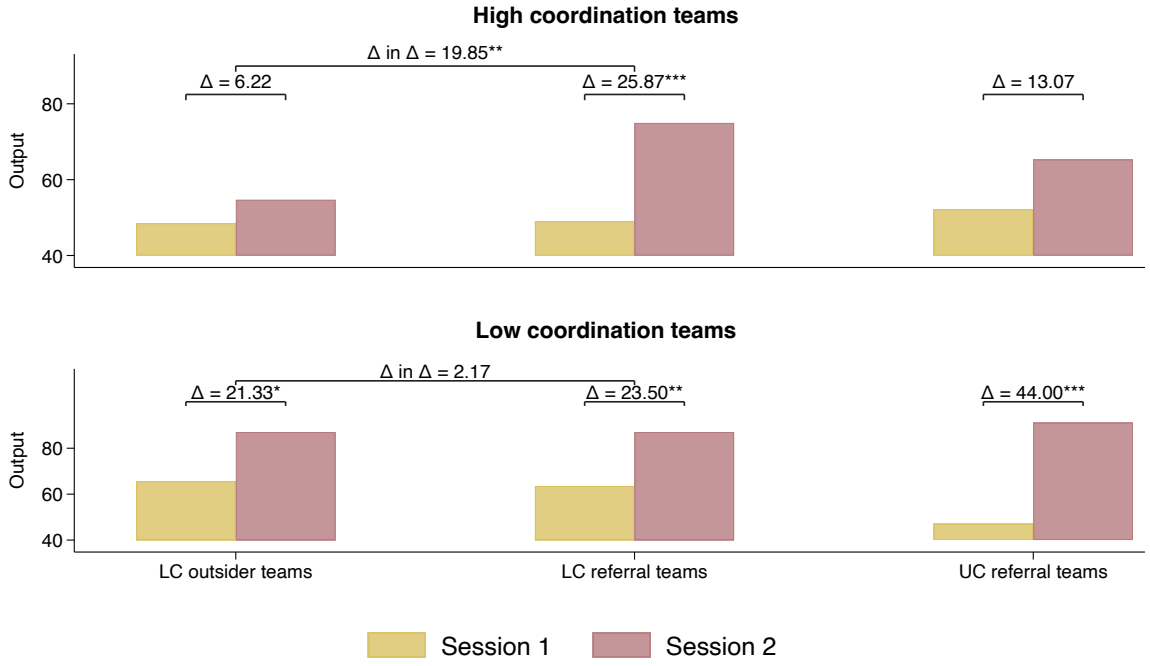
This figure plots the evolution of treatment effects on team-level output at the firm as a result of the firm experiment, relative to the timing of the first exit at the firm in the experiment (or post-experiment) period. The estimates come from separate regressions for each "relative week" which is the calendar week relative to the week of the first experiment exit, with fixed effects at the randomization strata level and standard errors clustered at the team level.

Table B6: Lab-in-field experiment: Compensating for lack of effort

	Worked hard to compensate for others		
	(1) All	(2) High coord	(3) Low coord
LC referral teams	0.006	0.163**	-0.336**
	(0.067)	(0.068)	(0.123)
UC referral teams	-0.073	0.124	-0.516***
	(0.084)	(0.092)	(0.104)
LC outsiders mean	0.268	0.155	0.513
R^2	0.02	0.06	0.21
Observations	312	215	97

This table reports responses to a question where workers are asked if they worked harder to compensate for others on their team than they would have expected. Standard errors are clustered at the team level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Figure B9: Summary of results from lab-in-field experiment



This figure plots the treatment effect on raw output in the lab-in-field experiment. The yellow bar plots the output from the first session while the red bar plots the output from the second session, when each team had three and four workers respectively. This figure shows that in high coordination tasks, LC outsider teams gain no detectable value from session one to two, from adding an additional worker. On the other hand, the gains are large and detectable in the referral teams. These differences are denoted through difference in difference estimates marked as appropriate. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table B7: Firm endline: beliefs

	Good at job		Hardworking		Good as team member	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.195	0.133	0.053	-0.154	0.227	0.444*
	(0.175)	(0.234)	(0.191)	(0.237)	(0.166)	(0.229)
Treatment × Lower caste		0.219		0.579		-0.617
		(0.356)		(0.596)		(0.376)
Control mean	-0.022	-0.022	-0.045	-0.045	-0.080	-0.080
R^2	0.09	0.09	0.07	0.08	0.10	0.11
Observations	208	208	204	204	205	205

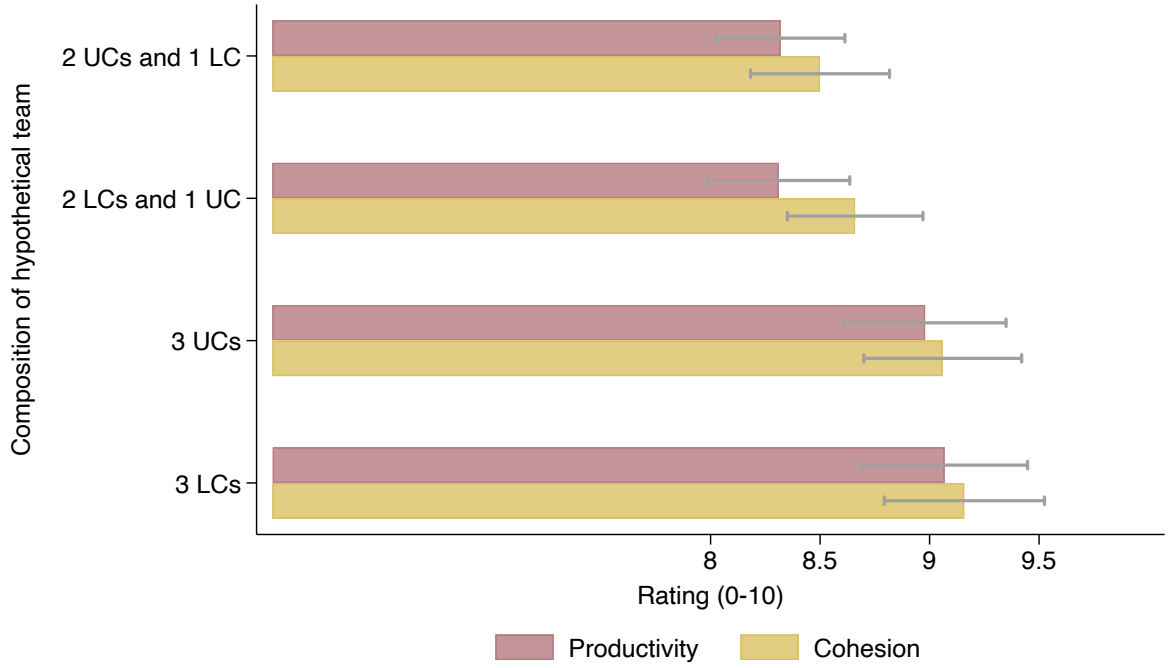
This table reports treatment effects on beliefs about out-groups, as reported in the endline survey at the firm. All standard errors are clustered at the team level, and the regressions inlcude fixed effects at the randomization strata level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table B8: Average tenure

	Tenure (months)	
	All workers	Current workers
Treatment	1.383***	1.741***
	(0.300)	(0.404)
Control mean	5.845	5.201
R^2	0.02	0.04
Observations	1365	796

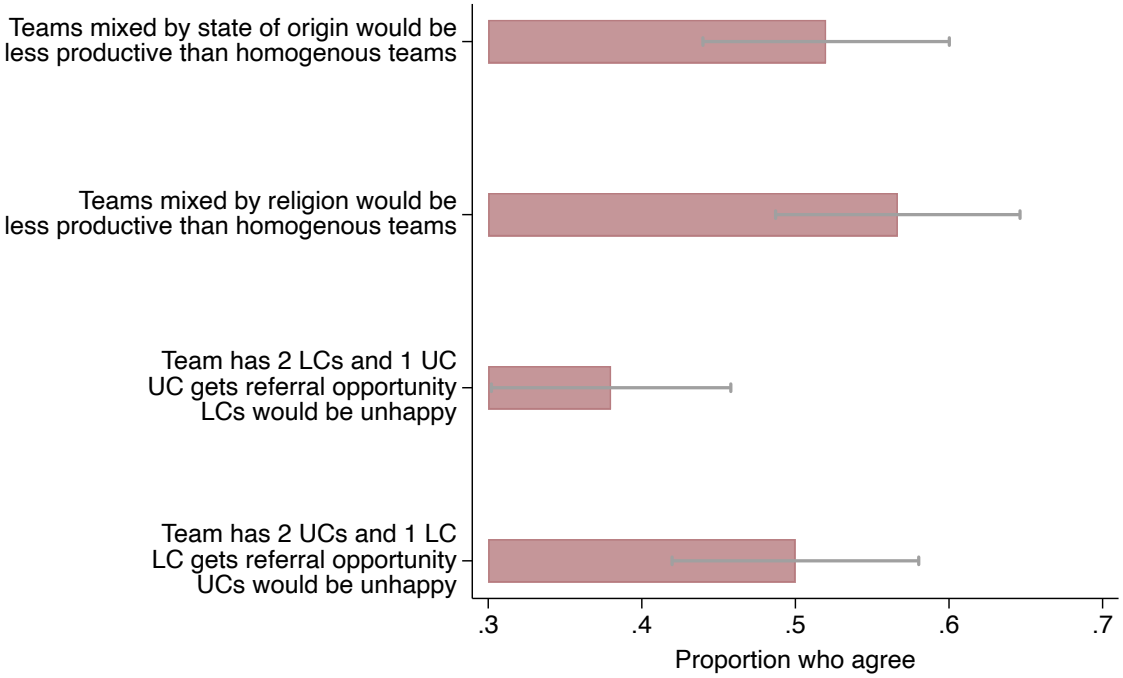
This table reports the tenure at the worker level at the firm, as of the endline. All standard errors are clustered at the team level, and the regressions inlcude fixed effects at the randomization strata level. $*$ $p <$ 0.1, $**$ $p < 0.05$, $***$ $p < 0.01$

Figure B10: Vignettes survey: expectations of productivity and cohesion vs. team composition



This figure summarizes responses from the vignettes survey on the expected productivity of hypothetical teams based on their composition.

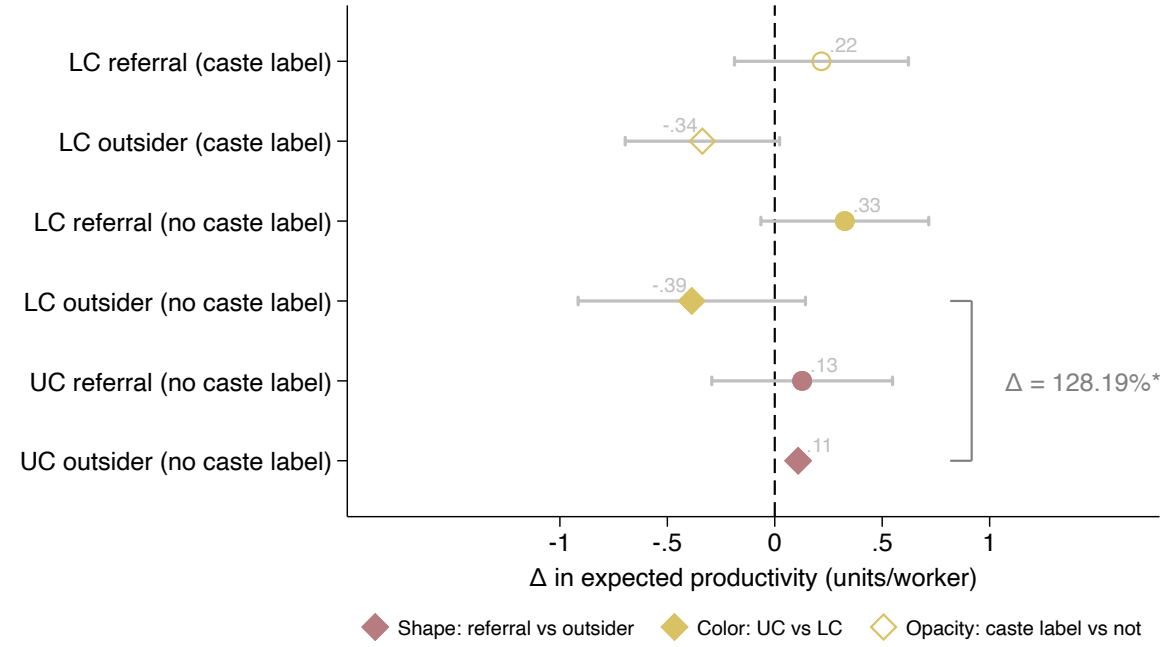
Figure B11: Vignettes survey: policy perceptions



This figure summarizes responses from the vignettes survey on agreement rates to a range of statements concerning team diversity, cohesion and productivity.

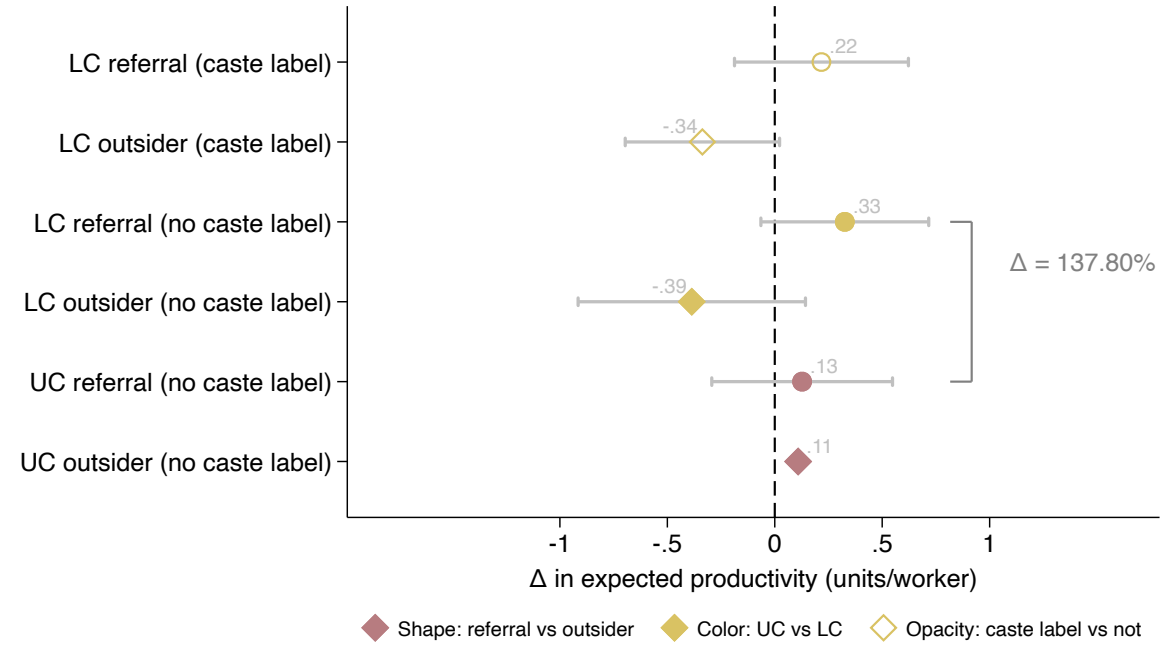
Appendix C: Vignettes survey

Figure C1: Vignettes survey



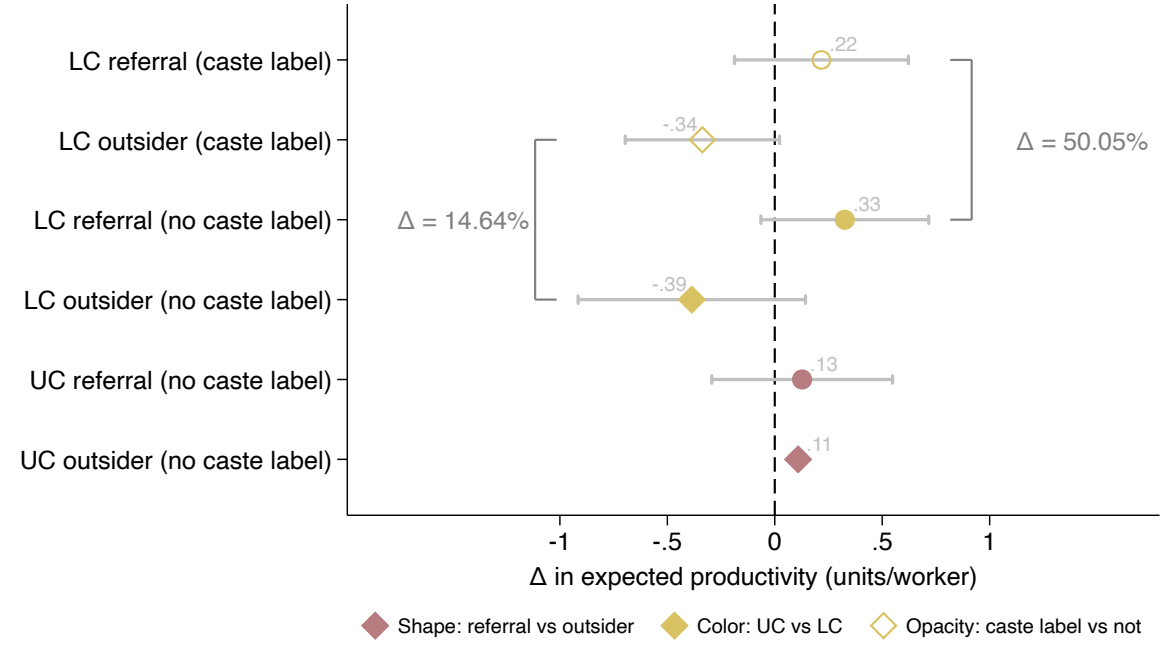
598 observations across 150 individuals; includes individual FE; CIs refer to pairwise comparisons with UC outsider estimate

Figure C2: Vignettes survey



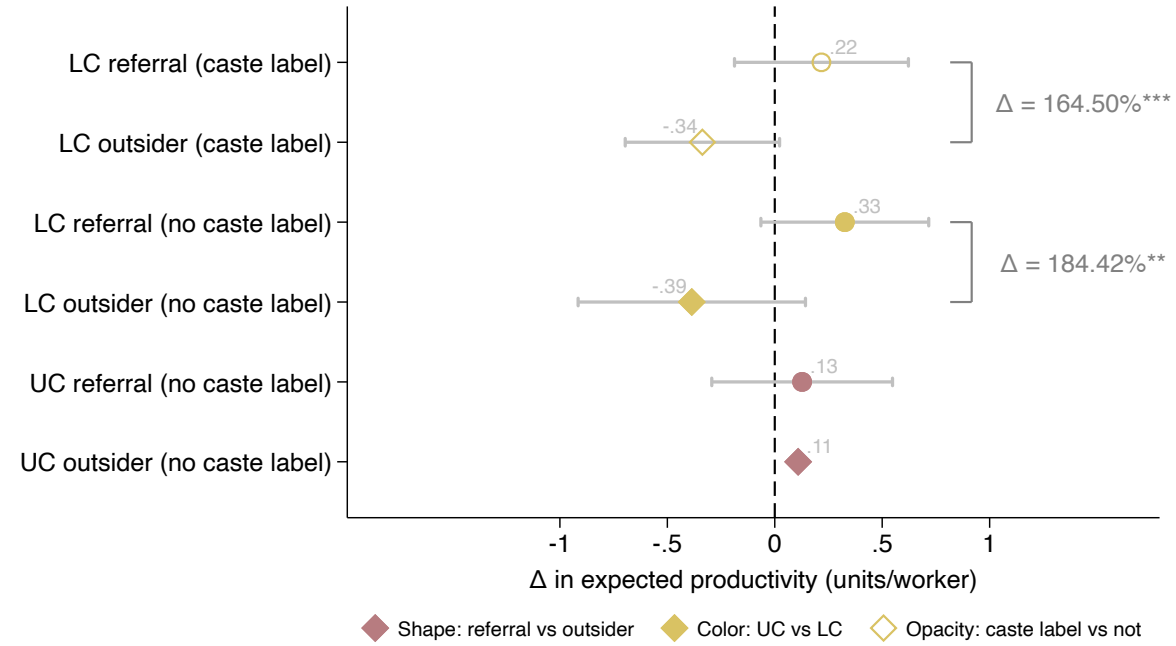
598 observations across 150 individuals; includes individual FE; CIs refer to pairwise comparisons with UC outsider estimate

Figure C3: Vignettes survey



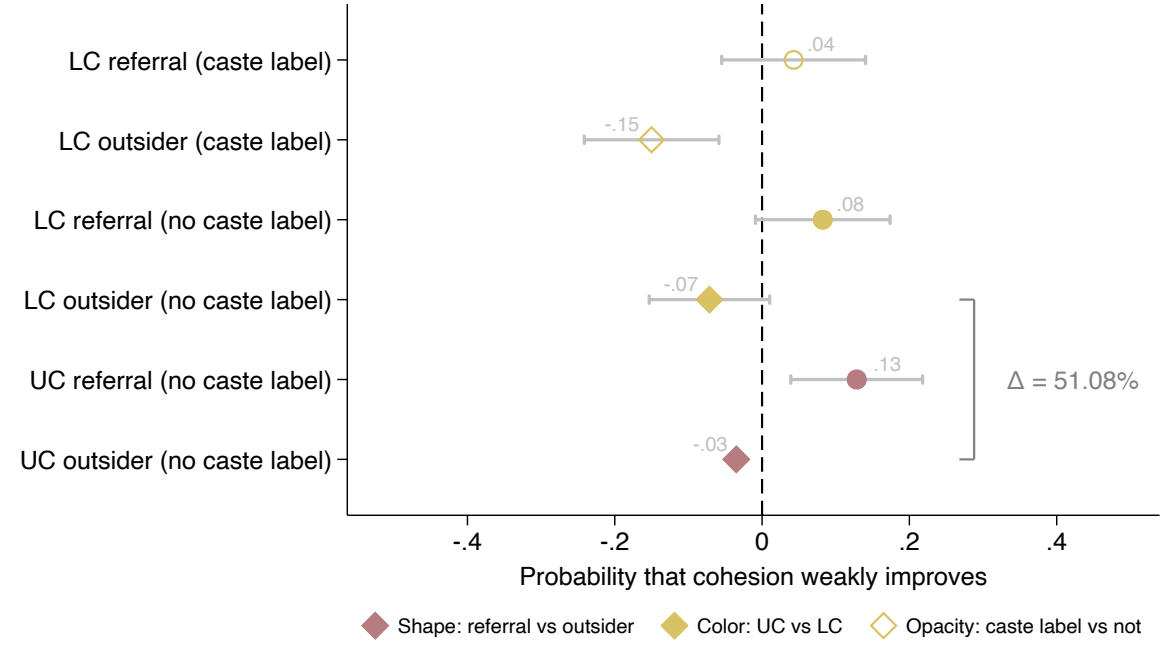
598 observations across 150 individuals; includes individual FE; CIs refer to pairwise comparisons with UC outsider estimate

Figure C4: Vignettes survey



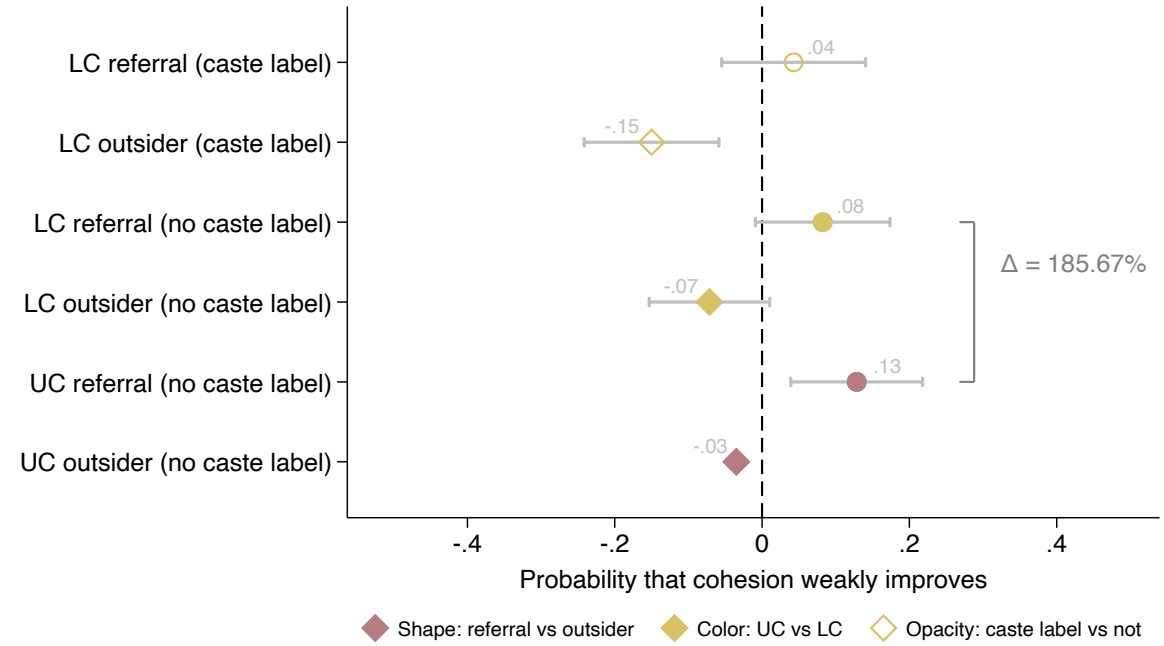
598 observations across 150 individuals; includes individual FE; CIs refer to pairwise comparisons with UC outsider estimate

Figure C5: Vignettes survey



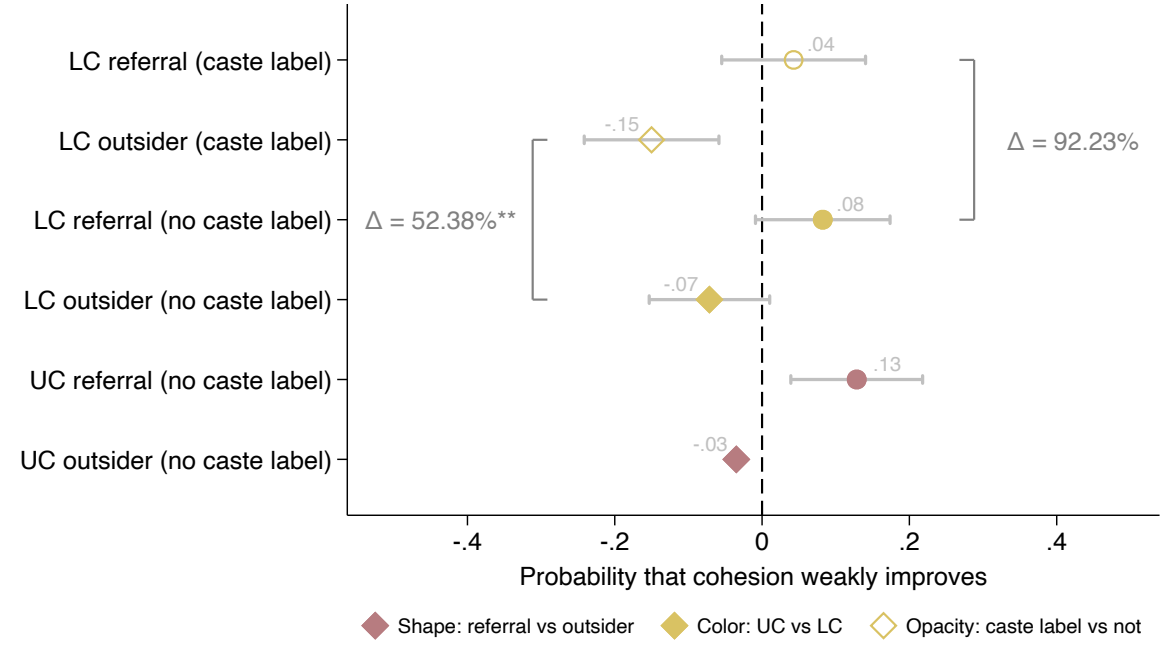
598 observations across 150 individuals; includes individual FE; CIs refer to pairwise comparisons with UC outsider estimate

Figure C6: Vignettes survey



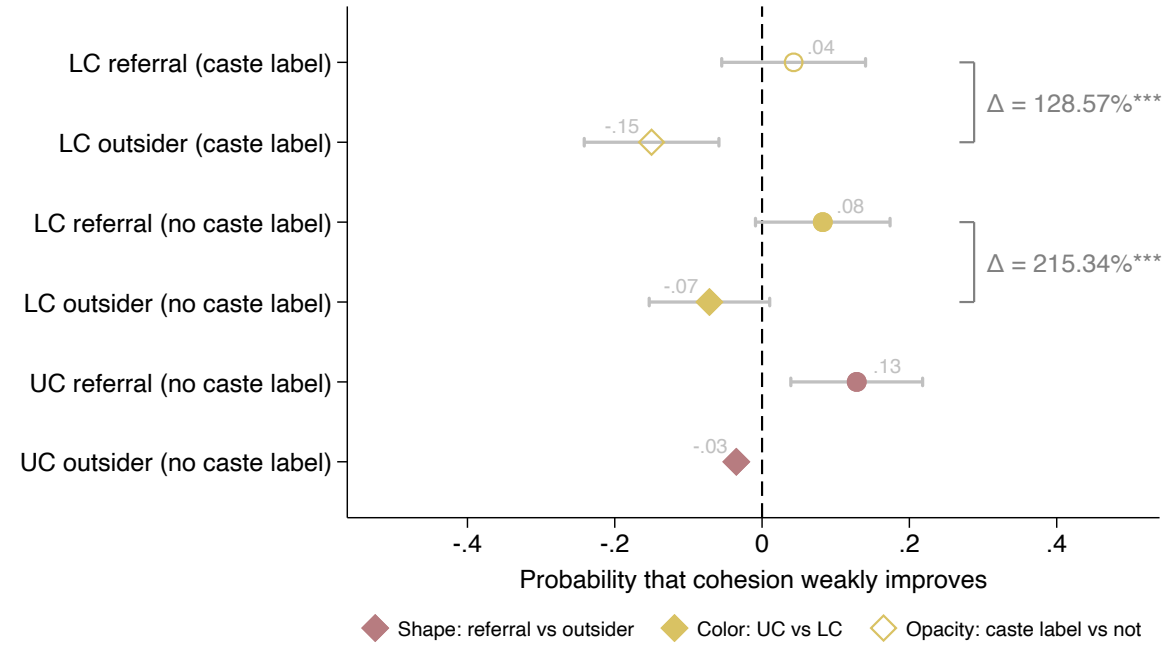
598 observations across 150 individuals; includes individual FE; CIs refer to pairwise comparisons with UC outsider estimate

Figure C7: Vignettes survey



598 observations across 150 individuals; includes individual FE; CIs refer to pairwise comparisons with UC outsider estimate

Figure C8: Vignettes survey



598 observations across 150 individuals; includes individual FE; CIs refer to pairwise comparisons with UC outsider estimate

Appendix D: Model

The firm must allocate referral opportunities to either lower caste or upper caste workers when vacancies arise, balancing team retention, cohesion, and productivity.

State variables: $S_t = (n_l^t, n_u^t, \nu_{\kappa,t}, \nu_{\tau,t}, \sigma^2_{\kappa,t}, \sigma^2_{\tau,t})$

- (n_l^t, n_u^t): number of LC and UC workers in period t respectively
- $(\nu_{\kappa,t}, \nu_{\tau,t})$: posterior means for beliefs about κ, τ
- $(\sigma^2_{\kappa,t}, \sigma^2_{\tau,t})$: posterior variances for beliefs about κ, τ

Control variable: X_l^t: number of vacancies filled by referrals from LC workers in period t

Let p_i^g indicate the probability that a worker i from group $g \in \{u, l\}$ will remain at the firm in any given period:

$$p_i^g \sim \text{Beta}(\eta^g, \phi^g)$$

For each worker i, whether they exit in period t is based on a Bernoulli random variable:

$$Q_i^{g,t} \sim \text{Bernoulli}(1 - p_i^g)$$

I assume here that lower castes have higher average retention, i.e. $\frac{\eta^l}{\eta^l+\phi^l} > \frac{\eta^u}{\eta^u+\phi^u} \implies \mathbb{E}[p_i^l] > \mathbb{E}[p_i^u]$.

The team composition evolves as follows:

$$n_g^{t+1} = n_g^t + X_g^t - Q_g^t.$$

I assume here that referrals bring others from the same group as themselves, i.e. $X_g \in g$.

The cohesion factor C^t depends on the team composition:

$$\begin{aligned} C^t &= 1 - \hat{\kappa_t} \cdot \rho_t - \hat{\tau_t} \cdot \Delta_t, \text{ where} \\ \rho_t &= \frac{n_l^t}{L} \cdot \frac{n_u^t}{L} = \lambda_t \cdot (1 - \lambda_t) \text{ and} \\ \Delta_t &= \mathbb{I}\{\lambda_t > \lambda_0\} \cdot (\lambda_t - \lambda_0). \end{aligned}$$

λ_t is the share of lower caste workers in the team in period t, λ_0 is the initial lower caste share of the team, $\hat{\kappa_t} > 0$ is the cost of cross-caste cohesion, and $\tau > 0$ is the extent of (taste-based) discrimination against LCs. Note that cross-caste cohesion has two components: the product of the caste-specific shares of the team, and an indicator term that accounts for whether there are more lower caste workers in period t than there were at the initial stage, scaled by the difference in the proportion of lower caste workers in period t and the initial stage.

In period t, the firm believes the scale parameter for the cohesoin cost is $\hat{\kappa_t}$ and the taste-based discrimination term is $\hat{\tau_t}$. Over-estimating κ and/or τ would result in the firm allocating fewer referrals to LCs than may be optimal. Under-estimating them would result in the firm allocating more.

After allocating the referral, the firm observes the actual cohesion in the team:

$$C^t = 1 - \kappa_t \cdot \rho_t - \tau_t \cdot \Delta_t + \epsilon_t,$$

where $\epsilon_t \sim N(0, \sigma^2)$.

The firm then updates its beliefs about κ_t and τ_t using Bayes' rule. The prior beliefs for both parameters are $\hat{\kappa_t} \sim N(\nu_{\kappa,t}, \sigma^2_{\kappa,t})$ and $\hat{\tau_t} \sim N(\nu_{\tau,t}, \sigma^2_{\tau,t})$. The likelihood of the observed cohesion is:

$$\tilde{C_t} | \kappa_t, \tau_t \sim N(1 - \nu_{\kappa,t} \cdot \rho_t - \nu_{\tau,t} \cdot \Delta_t, \sigma^2)$$

In each subsequent period $t+1$, the firm's posterior beliefs about the mean levels of κ and τ can be estimated as follows:

$$\nu_{\kappa,t+1} = \nu_{\kappa,t} + \frac{\sigma^2_{\kappa,t} \cdot \rho_t}{\sigma^2 + \sigma^2_{\kappa,t} \cdot \rho_t^2} \cdot (C_t - (1 - \nu_{\kappa,t} \cdot \rho_t - \nu_{\tau,t} \cdot \Delta_t))$$

$$\nu_{\tau,t+1} = \nu_{\tau,t} + \frac{\sigma^2_{\tau,t} \cdot \Delta_t}{\sigma^2 + \sigma^2_{\tau,t} \cdot \Delta_t^2} \cdot (C_t - (1 - \nu_{\kappa,t} \cdot \rho_t - \nu_{\tau,t} \cdot \Delta_t))$$

The posterior variances evolve as follows:

$$\sigma_{\kappa,t+1} = \left(\frac{1}{\sigma^2_{\kappa,t}} + \frac{1}{\sigma^2} \cdot \rho_t^2\right)^{-1}$$
$$\sigma_{\tau,t+1} = \left(\frac{1}{\sigma^2_{\tau,t}} + \frac{1}{\sigma^2} \cdot \Delta_t^2\right)^{-1}$$

Putting all of these together, the firm's output in period t is:

$$Y^t = [A \cdot C^t] \cdot (L^t)^\alpha \cdot (K^t)^\beta,$$

modeled on augmented Cobb-Douglus functional forms used commonly in the literature (Bloom, Sadun, and Van Reenen 2016; Bartelsman, Haltiwanger, and Scarpetta 2013; Hamilton, Nickerson, and Owan 2003b; Hall and Jones 1999; Ichniowski, Shaw, and Prennushi 1997).

Objective function

The firm maximizes the present value of expected profits:

$$V(S_t) = \max_{X_l^t, X_u^t} \left[\pi^t + \delta \cdot \mathbb{E}[V(S_{t+1})|S_t, X_l^t, X_u^t]\right],$$

where:

- $\pi^t = Y^t - w \cdot L^t - \mu \cdot X^t$ is the period profit.
- $\mu \cdot X^t$ is the cost of training attributed to the number of entries in period t.
- δ is the discount factor ($0 < \delta < 1$).

The Bellman equation for the firm's dynamic problem is:

$$V(S_t) = \max_{X_l^t, X_u^t} \left[[A \cdot C^t] \cdot (L^t)^\alpha \cdot (K^t)^\beta - w \cdot L^t - \mu \cdot Q^t + \delta \cdot \mathbb{E}[V(S_{t+1})]\right].$$

Constraints

1. Vacancy constraint:

$$X_l^t + X_u^t = Q^t.$$

2. Team size constraint:

$$n_l^{t+1} + n_u^{t+1} = n_l^t + n_u^t + X^t - Q^t.$$

3. Cohesion cost:

$$C^t = 1 - \kappa_t \cdot \rho_t - \tau_t \cdot \Delta_t + \epsilon_t,$$

where the perceived cohesion cost C^t used in the decision is:

$$\hat{C}^t = 1 - \nu_{\kappa,t} \cdot \rho_t - \nu_{\tau,t} \cdot \Delta_t$$

I estimate this dynamic model numerically using aggregate data from the firm experiment. Recall that the likelihood of a worker exiting is $Q_i^{g,t} \sim \text{Bernoulli}(1 - p_i^g)$, where $p_i^g \sim \text{Beta}(\eta^g, \phi^g)$. I use the baseline means of retention probabilities for each caste group in the firm to estimate the parameters of this distribution. For instance, if $M_{\text{pre}}^l, SD_{\text{pre}}^l$ represents the mean and standard deviation of the retention probabilities of the firm's LC workers at baseline, I estimate the paramaters as follows:

$$\eta^l = M_{\text{pre}}^l \cdot \left[\frac{M_{\text{pre}}^l(1 - M_{\text{pre}}^l)}{SD_{\text{pre}}^l} - 1\right]$$
$$\phi^l = (1 - M_{\text{pre}}^l) \cdot \left[\frac{M_{\text{pre}}^l(1 - M_{\text{pre}}^l)}{SD_{\text{pre}}^l} - 1\right]$$

Next, I use baseline surveys with supervisors to estimate initial measures of κ and τ as follows:

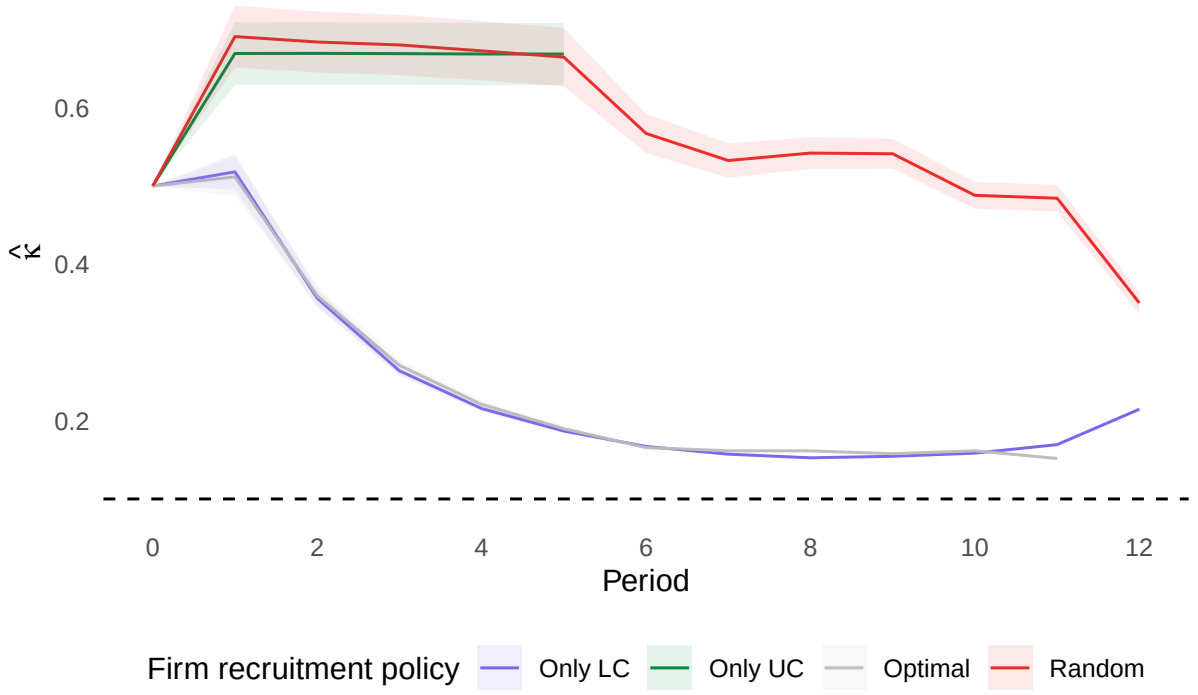
$$\nu_{\kappa,0} = 0.6; \qquad \sigma_{\kappa,0} = 0.2$$
$$\nu_{\tau,0} = 0.3; \qquad \sigma_{\tau,0} = 0.2$$

I then estimate the evolution of beliefs about κ and τ under four possible recruitment policies, simulated over a range of teams with varying initial shares of LCs at $t = 0$:

- **Only LC:** all vacancies are replaced through referrals to incumbent LCs, chosen randomly; *this policy replicates the "treatment" group from the firm experiment.*
- **Only UC:** all vacancies are replaced through referrals to incumbent UCs, chosen randomly.
- **Random:** all vacancies are replaced through referrals to a random incumbent worker; *this policy replicates the "control" group from the firm experiment.*
- **Optimal:** all vacancies are replaced through referrals allocated to the most optimal choice based on the discounted future value of making this choice with respect to the cohesion and training costs; *this policy replicates the schedule of referral allocations that a firm solving this dynamic problem will arrive at.*

For the purpose of this simulation exercise, I estimate that the true values of κ and τ are 0.1 and 0.01 respectively, and that each period the firm gets a signal of this true cohesion cost with noise such that $\sigma = 0.2$.

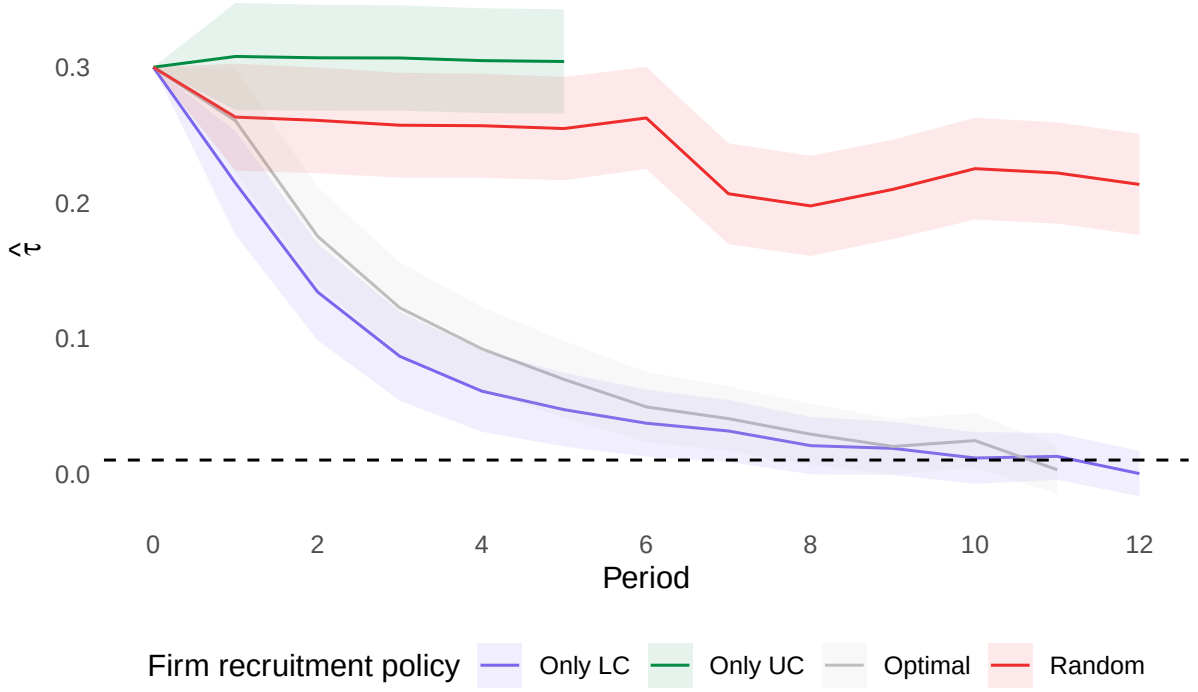
Figure D1: Evolution of the firm's beliefs ($\hat{\kappa}$)



Comparative statics

In this section, I present insights from a simplified static version of the model.

Figure D2: Evolution of the firm's beliefs ($\hat{\tau}$)



Consider a firm hiring workers from LCs and UCs. Let n^l and n^u denote the number of minority (lower-caste) and majority (upper-caste) workers, respectively, in the workforce. The firm fills Q_t vacancies each period via referrals from current employees: X_t^l referrals come from minority workers and X_t^u from majority workers, so that

$$Q_t = X_t^l + X_t^u.$$

New hires are from the same caste as their referrer (i.e., referrals are homophilous, bringing in same-group candidates). The workforce evolves as:

$$n_{t+1}^l = n_t^l + X_t^l - Q_t^l, \quad n_{t+1}^u = n_t^u + X_t^u - Q_t^u,$$

where exits denote those who leave (turnover). We assume that minority workers have higher retention on average than majority workers (i.e., the retention probabilities satisfy $p_l > p_u$).

The firm's *cohesion* in period t is captured by a factor $C_t \in [0, 1]$ that scales output. We model this as

$$C_t = 1 - \kappa\, \rho_t - \tau\, \Delta_t,$$

where

- $\rho_t = \dfrac{n_t^l}{L} \cdot \dfrac{n_t^u}{L}$ measures cross-group mixing (with $L_t = n_t^l + n_t^u$),

- $\Delta_t = \mathbb{I}\{\lambda_t > \lambda_0\} \cdot (\lambda_t - \lambda_0)$ captures any additional cohesion cost from increasing minority representation beyond a baseline.

Here, $\kappa > 0$ governs the cost of cross-group interactions (loss of cohesion from heterogeneity), and $\tau > 0$ represents a prejudice/discrimination cost that activates if minority presence exceeds the historical baseline.

The firm's output is given by an augmented Cobb–Douglas production function:

$$Y_t = [A \cdot C_t]\,(L_t)^{\alpha}(K_t)^{\beta},$$

where A is base productivity and α, β are the output elasticities of labor and capital, respectively. A lower cohesion factor C_t directly reduces effective productivity.

The per-period profit is

$$\pi_t = Y_t - w\,L_t - \mu\,Q_t,$$

where $w\,L_t$ is the wage bill and μ is the training cost per new hire. The term $\mu\,Q_t$ reflects the cost of replacing workers who leave, so μ captures the value of *stability*.

In a steady-state or long-run equilibrium, the firm chooses referral allocations (X^l, X^u) (or equivalently, the fraction $f = \frac{X^l}{Q}$ of referrals going to minority candidates) to maximize the intertemporal profit. At steady state, the workforce composition is stable, implying

$$X^l = Q^l = (1 - p_l)\,n^l, \quad X^u = Q^u = (1 - p_u)\,n^u.$$

Let $\lambda^* = \frac{n^{l*}}{L}$ denote the equilibrium minority fraction in the workforce. Then, the fraction of new hires that are minority is

$$f^* = \frac{X^l}{X^l + X^u} = \frac{(1 - p_l)\,n^{l*}}{(1 - p_l)\,n^{l*} + (1 - p_u)\,n^{u*}}.$$

Since $n^{l*} = \lambda^* L$ and $n^{u*} = (1 - \lambda^*)L$, we can simplify to:

$$f^* = \frac{(1 - p_l)\lambda^*}{(1 - p_l)\lambda^* + (1 - p_u)(1 - \lambda^*)}. \tag{1}$$

In equilibrium, the firm's optimal f^* must also maximize profit. By considering the marginal profit of allocating one more referral to a minority worker (instead of a majority worker), we derive an equilibrium first-order condition.

Marginal Output Effect: Adding a minority hire increases λ_t (minority share), which affects cohesion C_t.

$$C_t = 1 - \kappa \cdot \lambda_t(1 - \lambda_t) - \tau \cdot \mathbb{I}\{\lambda_t > \lambda_0\} \cdot (\lambda_t - \lambda_0),$$

so that

$$\frac{\partial C}{\partial \lambda} = \begin{cases} -\kappa \cdot (1 - 2\lambda_t) - \tau, & \text{if } \lambda_t > \lambda_0 \\ -\kappa \cdot (1 - 2\lambda_t), & \text{if } \lambda_t \leq \lambda_0 \end{cases}$$

Thus, the marginal effect on output is

$$\frac{\partial Y}{\partial \lambda} = \begin{cases} A\,L^{\alpha}K^{\beta}[-\kappa \cdot (1 - 2\lambda_t) - \tau], & \text{if } \lambda_t > \lambda_0 \\ A\,L^{\alpha}K^{\beta}[-\kappa \cdot (1 - 2\lambda_t)], & \text{if } \lambda_t \leq \lambda_0 \end{cases}$$

For $\lambda_t < 0.5$, since $1 - 2p_t > 0$, the effect on cohesion and output is negative for any $\kappa, \tau > 0$. For $\lambda_t > 0.5$, the effect on cohesion and output is positive as long as $\tau < \kappa \cdot (2\lambda_t - 1)$, i.e. the taste-based discrimination factor is not larger than the cross-group cohesion cost given the current team composition.[20]

Marginal Retention Effect: An LC hire, having a higher retention probability p_l compared to p_u, reduces future turnover. Specifically, the benefit is the reduction in future training costs, approximately $\mu\,(p_l - p_u)$ per hire.

Equilibrium First-Order Condition

At the optimum minority share λ^* (assuming an interior solution for $\lambda^* < 0.5$, which is the case for all but eight of the study teams over the course of the experiment at the firm), the marginal net benefit of shifting workforce composition is zero. Equating the marginal loss in output to the marginal gain from reduced training cost yields:

$$\text{FOC: } \begin{cases} A\,L^{\alpha}K^{\beta} \cdot [\kappa \cdot (1 - 2\lambda_t^*) + \tau] = \mu\,L\,(p_l - p_u), & \text{if } \lambda_t > \lambda_0 \\ A\,L^{\alpha}K^{\beta} \cdot [\kappa \cdot (1 - 2\lambda_t^*)] = \mu\,L\,(p_l - p_u), & \text{if } \lambda_t \leq \lambda_0 \end{cases} \tag{2}$$

Proposition 1

Proposition 1: *An increase in the cohesion cost parameter κ reduces the optimal fraction of minority referrals f^*, leading to a less diverse workforce and potentially lower long-run retention.* In equilibrium, $\frac{\partial p^*}{\partial \kappa} < 0$.

[20] Note that for the $\lambda_t > 0.5$ case, the effect on cohesion and output will always be positive when $\lambda_t \leq \lambda_0$, though that case is likely not binding given the relative value of λ_t.

Proof: The parameter κ amplifies the penalty for cross-group interactions. From Equation (2), a higher κ increases the marginal cohesion cost for any given λ. To restore balance, the firm must lower λ^* so that the term $(1 - 2\lambda^*)$ is smaller. This implies a reduction in λ^*, thereby reducing the minority share and shifting referrals toward the majority. Consequently, the firm avoids high cohesion costs but sacrifices the stability (retention) benefit from hiring minority workers.

Proposition 2

Proposition 2: *The taste-based discrimination parameter τ creates a threshold effect: if τ is large, the firm will cap minority representation at or below the historical baseline ($\lambda^* \approx \lambda_0$). Increases in τ (holding κ fixed) weaken the incentive to expand minority referrals beyond the initial share.*

Proof: For $\lambda^* > \lambda_0$, discrimination costs add an extra penalty per unit increase in λ (see Equation (2)). A larger τ raises the effective marginal cohesion cost. As a result, the optimal λ^* must be lower (or remain pinned at λ_0) to avoid excessive output penalties. Thus, if τ is sufficiently large, the firm will refrain from increasing minority referrals beyond λ_0.

Proposition 3

Proposition 3: *An increase in the training cost μ raises the value of worker retention, thereby increasing the optimal minority referral share f^*. In equilibrium, $\frac{\partial \lambda^*}{\partial \mu} > 0$.*

Proof: A higher training cost μ magnifies the benefit of reduced turnover. Since minority workers have higher retention ($p_l > p_u$), increasing μ raises the right-hand side of Equation (2). To balance the equation, λ^* must increase so that the marginal output loss (left-hand side) is higher. In the limiting case, if μ becomes very large, the firm is incentivized to pursue a more diverse (or even all-LC) workforce, provided that the cohesion costs do not outweigh the benefits from stability.

Proposition 4

Proposition 4: *The model admits at least one stable equilibrium workforce composition $(\lambda^*, 1 - \lambda^*)$ that balances cohesion costs with retention benefits. In equilibrium, firm output and profit are higher when the workforce is more stable (i.e., lower turnover) and cohesion costs are moderate. However, due to non-linear effects of diversity on cohesion, multiple equilibria may exist:*

- *A low-diversity equilibrium with a predominantly UC workforce (small λ^*) yields high cohesion but higher turnover.*

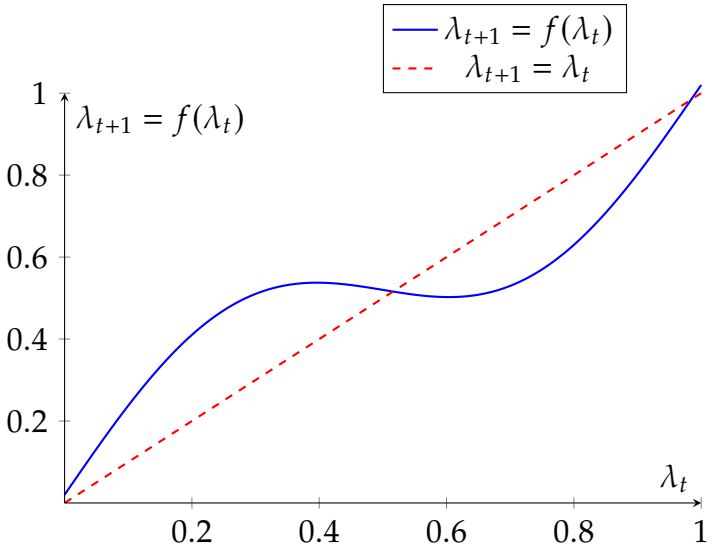


Figure D3: State transitions for LC share

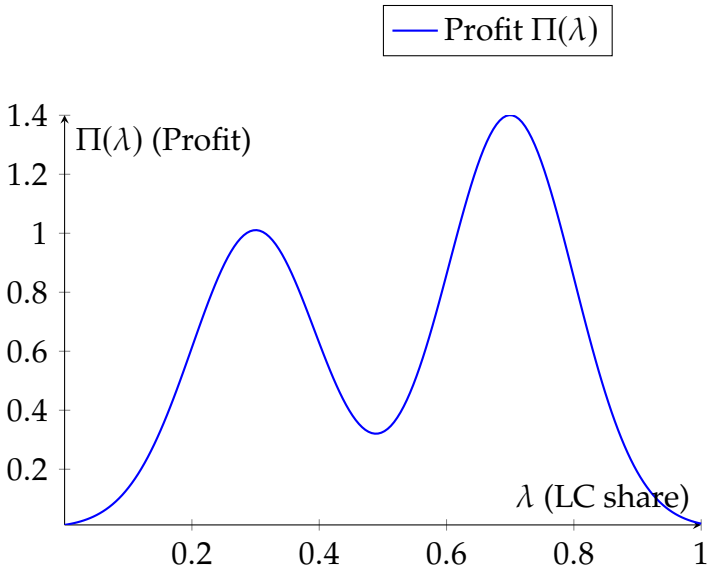


Figure D4: Profit relative to LC share

- *A high-diversity (or LC-dominated) equilibrium with large λ^* yields a homogeneous team (once LCs become the majority) and enjoys low turnover, potentially leading to higher long-run output.*

The referral process exhibits network externalities because incumbent workers tend to refer candidates similar to themselves. This homophily can sustain a given workforce composition. Moreover, if a policy intervention increases minority representation past a critical mass, the network effects may sustain a high-diversity equilibrium with improved stability and output. Path dependence and the existence of multiple local optima are inherent in this non-linear system.

In Figure D3, I plot the state transitions for the share of LC workers per period, demonstrating three equilibria: low LC share, intermediate LC share, and high LC share. The first and the third are stable, while the second is unstable. In Figure D4, I plot thr profit relative to the extent of LC shares. Note that while there is a symmetrical effect on cohesion for both the low-LC and high-LC equilibria, our assumptions on the model dynamics indicate that LCs have higher retention probabilities. This reduces training costs in the high-LC equilibrium, thus increasing profit relative to other equilibria.