Familiar Strangers:
Evidence from Referral-based Hiring Experiments in India

Kartik Srivastava*†

August 31, 2025

Job Market Paper

Updated very frequently. Click here for the latest version.

ABSTRACT. Labor markets in developing countries are shaped by high turnover, search frictions, and persistent exclusion of minority groups. While firms often rely on referral-based hiring to reduce search costs, this can systematically exclude those with worse networks. This paper asks whether reallocating referrals toward underrepresented workers can improve both equity and efficiency. In a manufacturing firm in India, I experimentally increased the share of referrals allocated to lower caste incumbents. The intervention raised their representation by 15 percentage points (62%) and reduced turnover by 4 percentage points (41%), driven by the longer tenure of lower caste entrants with weaker outside options. Output rose by 0.09 standard deviations, with no evidence of declines in team cohesion, in contrast to the literature on the adverse effects of mixed teams on workers' ability to coordinate. A lab-in-field experiment helps identify the mechanism: when lower caste workers are hired as outsiders, cohesion falls by 9% and output drops by 22% compared to when the same workers are hired through referrals, indicating that the mode of recruitment drives cohesion instead of the identity of the workers. Treated supervisors persist in allocating referrals to lower caste incumbent workers post-intervention, and lower caste referral candidates see large gains in employment outcomes – without diminishing the labor market prospects of their upper caste counterparts, who are able to find similar jobs elsewhere. These results highlight the untapped efficiency and equity gains from network-conscious hiring of underrepresented groups in frictional labor markets.

*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, Jeff Liebman, Amanda Pallais and MR Sharan. I thank the participants of the Development, Labor, and Economic and Social Policy workshops at Harvard University, as well as workshop participants at Symposium on Economic Experiments in Developing Countries, HEC Paris, University of Exeter, Warwick University, European Association of Young Economists Annual Meeting, Scotland and Northern England Conference in Applied Microeconomics, and Queen Mary University of London. 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

About 125 million people in developing countries were unemployed in 2024. Yet, more than 70% of employers reported difficulty finding suitable candidates (ILO, 2024; Manpower Group, 2024). Labor markets in these settings are marked by exceptionally high turnover: job switching among wage workers is about three times higher in India and Ecuador than in the United States or Sweden, and churn in these settings tends to be lateral and not suggestive of progression on the job ladder (Donovan et al., 2023). Meanwhile, even well-paying firms struggle to hire and retain workers, despite widespread unemployment and limited upward mobility. In this context, many empoyers rely on referral-based hiring – where incumbent workers are asked to recommend candidates from their social networks to fill a vacancy. This practice is especially common in frictional labor markets, where search and screening costs are high and moral hazard is a risk (Breza and Kaur, 2025; Chandrasekhar et al., 2020; ILO, 2018).

Hiring through referrals, though, tends to favor well-connected job-switchers. As a result, workers with weaker or poorly connected social networks are more likely to be left trapped in low-earning self-employment or long unemployment spells (Beaman et al., 2018). This dynamic is particularly damaging in contexts where job access requires migration or cross-sector transitions. Firms may become trapped in a low search cost, high turnover equilibrium, systematically excluding minority workers with a high willingness to accept these jobs. Such referral patterns can distort labor allocation and reinforce segmentation, pointing to a market failure.

This paper studies the trade-off firms face between improving worker retention and managing the potential coordination costs of more demographically diverse teams. Concentrating referrals among frequent referrers can lower marginal search costs, but may attract candidates with strong outside options, increasing churn (Calvo-Armengol and Jackson, 2004; Bolte et al., 2021). Conversely, firms might worry that raising the share of underrepresented workers could trigger inter-group conflict and hamper workplace cohesion (Ghosh, 2023; Hjort, 2014; Lazear, 2000). I examine these competing forces in entry-level employment in the context of India's caste system – an enduring institution that continues to shape occupational mobility and labor market access (Cassan et al., 2022; Lowe, 2021; Munshi, 2019; Oh, 2023).

I examine this question using four complementary approaches. First, I conduct a large-scale field experiment at a manufacturing firm in New Delhi, where referral allocation was experimentally varied over six months among nearly 800 workers. Second, I implement a lab-in-field extension that isolates mechanisms underlying the firm-level findings. Third, I run a vignette-based survey experiment to assess external validity and worker perceptions of different recruitment policies. Finally, I develop a dynamic Bayesian learning model to study how supervisors update beliefs about cohesion costs and retention, and how these beliefs shape referral allocations at the firm over time.

Access to referral opportunities is both a gatekeeper to employment and a mechanism through which existing inequalities in hiring are perpetuated. At baseline, over 90% of frontline workers at the firm had been hired through referrals from incumbent employees.

Yet, lower caste workers – who constituted 21% of the workforce – received only 12% of referral opportunities. Typically, team supervisors selected which incumbent received the right to refer when a vacancy arose. During the experiment, I altered this process: in treatment teams, referral opportunities were randomly allocated exclusively to lower caste incumbents; in control teams, referrals were randomly assigned regardless of the caste of the incumbent.

Compliance with the assigned referral rules was high: 92% on average at the team-by-month level. Over six months, the share of lower caste workers in treatment teams increased by 15 percentage points—a 62% rise ($p = 0.01$). Retention also improved: monthly churn fell by 4 percentage points (41%) relative to control teams, driven primarily by fewer exits among newly hired lower caste workers. These teams experienced 0.6 fewer exits and 0.7 fewer entries per month on average. In contrast, control teams continued to cycle through new hires at their baseline rate. Baseline survey data indicate that lower caste referrals were 65% more likely to be unemployed and disproportionately based in rural areas compared to their upper caste counterparts, suggesting that their weaker outside options contributed to longer tenures at the firm.

The potential impact of increased lower caste representation on productivity is theoretically ambiguous. One possible concern is human capital and inherent ability: lower caste workers may, on average, have lower levels of formal education or training, despite steady gains in recent decades (Hnatkovska et al., 2013).[1] A second concern arises from team dynamics: more demographically mixed teams may suffer from weaker cohesion, potentially reducing productivity. A growing literature documents short-run productivity losses in diverse teams, often due to increased coordination costs or inter-group frictions (Hjort, 2014; Ghosh, 2023).[2]

At the same time, there are clear reasons to expect gains. Lower caste recruits tend to have weaker outside options, and thus are expected to stay longer, reducing costly turnover. Moreover, diversity may generate complementarities that raise performance over time. In my setting, I find that the net effect is unambiguously positive.

Over the six-month period, team-level productivity rose by 5% (0.09 standard deviations, $p = 0.03$), with particularly pronounced gains in tasks requiring close coordination among teammates. This pattern aligns with prior evidence that turnover is especially disruptive in interdependent work environments (Hamilton et al., 2003a; Huckman and Staats, 2013). Importantly, I find no evidence of the short-run productivity declines often observed in diverse teams. Output in treatment and control teams was similar at the outset, but diverged over time as treatment teams stabilized and settled into more efficient work routines. Survey measures of team cohesion and job satisfaction also show no statistically

[1] Since these are entry-level jobs, neither prior experience nor inherent ability is expected to play a major role in hiring. Indeed, baseline surveys show that neither supervisors nor workers perceive lower caste workers as intrinsically less capable or productive.

[2] Hjort (2014) finds short-term productivity declines in ethnically diverse teams in a Kenyan manufacturing firm. Ghosh (2023) shows similar effects when varying religious composition in an Indian plant, particularly for tasks requiring cooperation. In my study, treatment assignment is stratified by the coordination intensity of tasks to account for such variation.

significant differences at endline.[3]

These findings raise two natural questions. First, if directing referrals to lower caste workers improves retention and productivity, why were supervisors not doing so already? Baseline surveys reveal that supervisors hold priors consistent with concerns in the diversity literature: they anticipate that hiring lower caste workers through referrals to incumbent lower caste workers could reduce team cohesion and productivity. These fears appear to shape their referral allocations.[4]

The second question is why, despite these concerns, cohesion costs did not materialize in the experiment. To investigate, I conducted a vignette-based survey with workers drawn from a similar population. Respondents were asked to evaluate how team cohesion and productivity would be affected by the addition of a new member under different recruitment scenarios. The results are striking: lower caste recruits hired as "strangers" or outsiders, i.e. with no existing links to team members, are viewed as disruptive. In contrast, when the same caste-identifying entrants are introduced via referrals from incumbents, they are seen as neutral with respect to cohesion. These perceptions closely mirror the firm-level results: reallocating referrals to underrepresented incumbents did not trigger the cohesion losses commonly associated with diversity-focused hiring interventions.[5]

To further test these patterns in a more controlled setting, I conducted an incentivized lab-in-field experiment with nearly 500 participants, again drawn from the same population as the firm's workforce as in the vignette-based survey. Held across 12 workshops, this setting allows for tight control over task structure, recruitment conditions, and data collection – closely replicating the firm's production environment. Each team began with three members (one lower caste, two upper caste) and completed two production rounds, separated by a recruitment phase that increased team size to four. I randomized teams into one of three recruitment modes: (1) referral by the lower caste incumbent; (2) referral by an upper caste incumbent; and (3) direct hiring of a lower caste "stranger" without incumbent links. The first two mirror the treatment and control arms of the firm experiment, while the third mimics diversity-focused hiring policies that ignore existing team ties.

The results from the lab-in-field experiment echo those from the firm and shed light on why cohesion costs do not emerge. All teams begin with similar output levels and experience the same nominal increase in diversity. Selection is held constant across arms by recruiting candidates similarly prior to allocating them to teams as either referrals or

[3] Although self-reported data have limitations, these nulls suggest that no large declines in cohesion or morale were triggered. I supplement these measures with revealed-preference data on cooperation in the lab-in-field extension.

[4] Other explanations are possible. Supervisors may hold preferences against mixed teams that outweigh expected productivity gains, or they may be risk-averse, overweighting potential short-term disruptions even when long-term returns are positive. This would be consistent with a principal-agent problem between supervisors and firm leadership.

[5] There are parallels to this pattern in other contexts where agents' ability to choose improves cohesion. Experimental evidence suggests that endogenously chosen policies are more likely to foster cooperation and acceptance than those imposed from above (Dal Bó et al., 2010). In the context of team formation, Anders and Pallais (2025) show that teams formed discretionarily through friend networks exhibit stronger cohesion than those created exogenously.

"strangers", following Pallais and Sands (2016). Yet recruitment mode matters: in teams where the fourth worker is labeled a lower caste stranger, there is no significant gain in output. In contrast, when a similar worker is framed as a referral from an incumbent, team output rises by 22% ($p < 0.01$). As with the firm data, these gains are concentrated in tasks requiring high coordination, reinforcing the idea that networked integration, rather than diversity per se, drives performance improvements.

Unlike the firm setting, the lab-in-field experiment enables revealed-preference measurement of cohesion, using adaptations of behavioral cooperation games (Schofield, 2014; Breza et al., 2017). These tasks capture how well randomly matched pairs of workers within teams cooperate, suggesting higher workplace cohesion. Cohesion is 9% higher in lower caste referral teams than in lower caste stranger teams ($p = 0.02$), a gap that appears both in incumbent-only pairings and in pairings involving the new hire. This suggests that the recruitment mode shapes overall team cohesion, not just the dyad between referrer and recruit. Disaggregating by caste pairing further reveals that the treatment effect is driven by stronger cooperation in cross-caste pairs; no significant effects are observed among same-caste pairs.

The referral-based hiring intervention had lasting effects on firm behavior. More than five months after the experiment concluded, treatment teams continued to exhibit higher retention and productivity, and supervisors persisted in allocating referrals disproportionately to lower caste workers. This suggests that supervisors in the treatment group updated their beliefs through exposure to the enforced policy. Notably, even after the intervention ended, the share of referrals going to lower caste workers in treatment teams remained above their (now increased) share of the team's composition. In contrast, supervisors in control teams largely reverted to pre-intervention referral patterns.

To interpret these trends, I develop a dynamic model of referral allocation in which supervisors begin with prior beliefs about the cohesion costs of cross-caste hires and may also exhibit taste-based discrimination. Calibrated with baseline survey data and incorporating Bayesian updating under different referral policies, the model shows that pessimistic priors about cohesion can prevent supervisors from experimenting with lower caste referrals – even when doing so would improve retention and team stability in the long run.

This framework parallels the literature on strategic experimentation and multi-armed bandit problems (Gittins and Jones, 1979; Bolton and Harris, 1999; Aghion et al., 1991; Exley and Kessler, 2024; Li et al., 2025), where agents facing uncertainty avoid exploring potentially better options due to risk aversion or short-term incentives. In my setting, mandatory experimentation removes these frictions: by making exploration costless and enforced, the intervention allows treatment supervisors to update their beliefs and adopt more inclusive referral strategies that persist beyond the experimental period.[6]

In order to be able to assess downstream effects on job seekers, I collected detailed information on the social networks of incumbent workers at baseline, prior to treatment

[6] Li et al. (2025) study hiring as a form of exploration, emphasizing the trade-off between exploiting known talent and discovering underrepresented workers who may be equally or more productive. Their model closely parallels mine.

assignment. This allows me to compare outcomes for both upper and lower caste referral candidates in treatment versus control teams, focusing on those named by incumbents as potential referrals. Among lower caste candidates linked to treatment teams, I find a 52% increase in the number of job offers received relative to their control counterparts. In contrast, there is no corresponding effect for upper caste candidates. These improved job prospects translate into a 17% higher likelihood of securing salaried employment and a 29% reduction in the duration of the most recent unemployment spell. Crucially, these gains for lower caste candidates do not come at the expense of upper caste workers. I find no evidence of displacement: upper caste candidates in treatment teams secure job opportunities at rates comparable to those in the control group.

In this setting – as in many low- and middle-income countries – turnover tends to reflect lateral movement among marginal jobs rather than upward mobility (Donovan et al., 2023).[7] In original analysis of India's nationally representative 2018–19 Periodic Labor Force Survey, I find a strikingly bimodal distribution of outside options. Lower castes make job-to-job transitions 25% less often than upper castes, are 46% more likely to be in rural areas, and 13% more likely to be unemployed, controlling for a range of other demographic characteristics like age, gender, and region of birth.[8] These aggregate patterns suggest not only frictional labor markets and spatial misallocation, but also deeper segmentation, driven in part by differential access to referral networks that mediate entry into better-paying jobs.[9]

These qualitative findings are reflected in my empirical results. First, at baseline, lower caste referral candidates are significantly more likely to be unemployed or located in rural areas, consistent with weaker outside options. Second, endline survey data show that the intervention substantially improved labor market outcomes for lower caste candidates, with negligible costs to their upper caste counterparts, who face better baseline prospects. This is consistent with canonical models of occupational choice, where outside options and labor market trajectories are shaped by initial wealth and opportunity distributions (Banerjee and Newman, 1993).[10]

This paper contributes to a broad literature on labor market dynamics, social networks,

[7] Donovan et al. (2023) show that workers who transition to higher-wage jobs in developing countries are 39 percentage points more likely to exit employment or revert to lower-wage work within a quarter. More broadly, developing economies exhibit job-to-job transition rates twice as high and occupational switching rates four times as high as those in richer countries, with little associated wage progression.

[8] Using eight high-frequency employment panels (2017-19), Abraham and Kesar (2025) show that only 6.7% of workers ever reach stable formal salaried jobs, while over 70% cycle among precarious informal trajectories; access to each trajectory is "highly dominated by the privileged caste groups," whereas lower-caste workers are over-represented in the most volatile and lowest-paid paths.

[9] Similar asymmetries appear elsewhere. Fadlon et al. (2025) find that in the U.S., Black workers experience penalties in both upward and downward job transitions relative to White workers.

[10] These estimates are, however, partial-equilibrium. They assume other firms and market wages remain unchanged. In a general-equilibrium setting, widespread adoption of the policy would gradually raise the outside options of lower-caste workers and could compress the retention premium that drives these results. I discuss in the structural model below why such wage adjustments are likely to be gradual in fragmented manufacturing labour markets, and why persistent differences in network access mean the mechanism should retain much of its force even after partial wage equalisation.

and workplace diversity – particularly in contexts where segmentation and network access shape hiring. It builds on foundational models of job search and information frictions (Jovanovic, 1979; Montgomery, 1991) and a growing body of work on the costs of high turnover in firms in developing economies (Adhvaryu et al., 2023; Donovan et al., 2023; Halvorsen, 2021; Shiferaw et al., 2020; Skelton et al., 2019; Tybout, 2000). Closely related is the literature on social networks in recruitment, which spans both developed and developing country contexts. Studies in high-income settings typically use matched employer–employee data to compare outcomes for referred versus non-referred workers (Brown et al., 2016; Burks et al., 2015; Dustmann et al., 2016; Kugler, 2003), while work in developing countries emphasizes referrals as a mechanism for mitigating adverse selection and moral hazard (Heath, 2018; Pallais and Sands, 2016) or as a lens into referrer behavior and candidate selection (Beaman and Magruder, 2012; Beaman et al., 2018). This work contributes to this literature by focusing on an upstream but understudied aspect: how referral opportunities are allocated among incumbent workers, and how this allocation affects team composition, retention, and productivity.

This paper also relates to a classic line of inquiry in labor economics on spatial and informational frictions, particularly the exclusion of disadvantaged groups from jobs they are otherwise willing to accept (Chetty et al., 2014; Kain, 1968; Moretti, 2012). The empirical findings demonstrate how limited network access can sustain misallocation, even when willingness to work is high.

Methodologically, this work is distinctive in combining intra-firm variation in referral policies with evidence from a real production environment. Prior research often relies on administrative data – where referral links are inferred as opposed to identified – or on stylized experiments where researchers act as employers. These approaches may not be able to fully capture the relational dynamics between firms and incumbents that shape referrals and workplace integration. By embedding the intervention within an existing firm and complementing it with lab-in-field experiments, this paper offers novel evidence on how referral structures evolve and how hiring frictions can be relaxed (Bandiera et al., 2011; Hoffman and Stanton, 2024).

The findings also speak to literatures on workplace cohesion and the role of management in shaping productivity (Lazear, 2000; Mas and Moretti, 2009; Roethlisberger and Dickinson, 1939; Bandiera et al., 2005; Bandiera et al., 2009; Fershtman and Gneezy, 2001; Prat, 2002). A related strand examines how team diversity affects trust and output, often identifying short-run costs from increased coordination burdens or latent biases (Alesina and La Ferrara, 2005; Ottaviano and Peri, 2006; Ghosh, 2023; Hjort, 2014). This paper demonstrates how a network-based recruitment strategy – specifically, referrals from incumbent workers – can reduce turnover and enhance performance without incurring the cohesion losses that often accompany diversity interventions.

Beyond economics, the paper connects to insights from sociology and social psychology on intergroup contact, individuation, and stereotype updating. These literatures emphasize that individuals may revise beliefs about members of other groups only when those interactions are broad-based and repeated (Fiske and Neuberg, 1990; Granovetter, 1973; Brewer, 1999). As Chakraborty et al. (2024) show, exposure to many outgroup

members is more effective at reshaping generalized beliefs, while deep ties foster long-term collaboration. In my setting, referral-based hiring allows for positive associations to extend from individual lower-caste workers to their direct social contacts – facilitating integration even if broader out-group attitudes remain unchanged.

The paper also has theoretical roots in models of networked labor markets, including those of Calvo-Armengol and Jackson (2004) and Bolte et al. (2021), who show how network structures shape employment allocations. My experiment tests these predictions in a real-world setting, linking referral networks to retention, cohesion, and productivity outcomes.

Finally, this work carries direct policy relevance. Referral-based hiring is widespread across the developing world (Afridi et al., 2021; Afridi et al., 2020; Witte, 2021), but often reinforces homophily and exclusion. Networks can perpetuate inequality by concentrating access to jobs among connected individuals, leading to inefficient equilibria with high turnover and underutilized labor pools. These patterns mirror findings across gender (Beaman et al., 2018), race (Arrow and Borzekowski, 2004), and migrant status (Munshi, 2003; Patacchini and Zenou, 2012).[11] By reallocating referral opportunities toward under-represented groups, firms can unlock latent talent. My findings suggest that embedding diversity within existing social ties can yield both equity and efficiency gains.

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 extension 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 industrial 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

[11] For example, Munshi (2003) shows that among Mexican migrants in the U.S., the size of a worker's network determines both access to jobs and the quality of those jobs.

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.

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 et al., 2023; Beaman et al., 2018).[12]

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

[12]Referrals are used commonly in other (non-labor market) contexts in developing countries as well. Bryan et al. (2015) show that peer referrals improve the enforcement of loan repayments in the microcredit context.

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 (World Bank, 2013).

Incumbent workers also value referrals, either by virtue of the status utility they confer or in a more general sense as a non-wage amenity. Friebel et al. (2023) show in an experiment that non-referral workers are more likely to stay at a firm when more referrals are hired, owing to the fact that they derive utility from being able to have a voice in choosing their coworkers. In my context, as well, I find that workers have a substantial willingness to pay for being able to allocate referrals: 74% are willing to trade off between 10 - 30 additional minutes of their regular lunch break in order to be able to allocate a referral.

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 and the implications of diversity

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 remaining are classified into the "upper caste" category.[13]

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 et al., 2013; Asher et al., 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.[14] As a result, SCs often remain confined to low-paying, menial jobs

[13]This is a coarse classification, which necessarily introduces some errors of inclusion and exclusion. A granular version of this analysis would focus more squarely at the *jati* level. However, given the saturation of caste groups represented in my sample who happen to have migrated from a limited number of sending regions, this coarse classification is not a first order concern. Moreover, an analysis at the *jati* level would generally not be powered at conventional levels.

[14]Caste-based networks dominate the dissemination of both information about job opportunities and

with limited upward mobility, perpetuating economic disparities and social marginalization. Network-based job search has previously been used to explain persistent caste-based differences in occupational sorting and wage differences in urban areas (Banerjee and Knight, 1985). 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 B.1). Even when they are able to migrate to cities, they are concentrated in low-premium industries.

Consistent with the literature on the contact theory, there is evidence that exposure to caste based out-groups changes beliefs about these out-groups. In my setting, we can take the set of existing lower caste workers in teams as coming from some pre-determined selected population wherein through contact theory workers are not biased towards them in particular (Finseraas et al., 2019; Lowe, 2021). Recent work stresses that it is broad exposure to a large number of out-group members that leads priors to be updated, as opposed to deep exposure to a smaller number (Chakraborty et al., 2024). In line with this, in my setting it is reasonable to expect that incumbent upper caste workers may still hold negative beliefs about lower caste workers at large.

A vast literature suggests that when firms are engaged in team-based production, diverse teams tend to be less productive in the short run owing to higher costs of coordination and cooperation among workers. This is shown in the context of religious diversity in India (Ghosh, 2023), where teams made of diverse Hindu and Muslim workers are less cohesive owing primarily to misplaced beliefs about the productivity and shirking of out-groups. There is similar evidence from ethnic groups in Kenya (Hjort, 2014), where there is horizontal and vertical discrimination within teams. In other contexts, there are more direct coordination costs from diversity, such as when workers who speak different languages are asked to coordinate on tasks (Dale-Olsen and Finseraas, 2020; Anisfeld and Pandiloski, 2024). Here too, there are declines in output in mixed teams that ameliorate over time. Álvarez Pereira et al. (2024) show that diversity is more beneficial in larger teams.

A related literature suggests that there is a risk of substantial backlash against affirmative action policies in the context of caste in India. In particular, Jodhka and Newman (2007) show that managers at firms express a distaste towards affirmative action policies, and use that as an explanation for the lack of access for lower caste workers in the formal private sector in India. In other contexts as well, affirmative action policies have been shown to harden identities and polarization, such as in the case of Brazil (Pérez Cervera, 2023). Xu (2021) shows in an audit study that putatively Black applicants who are shown to have received scholarships labeled as diversity (or without labels) become less likely to receive call backs from prospective employers, while putatively White applicants face no such penalties.

provide pathways that drive migration even in the absence of direct job opportunities, through social insurance that supports periods of unemployment (Mohanty, 2024).

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 B.1). 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 supervisor 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 et al. (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) reports a rate of 10% in the US context. [15]

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 B.11).[16] 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%).

[15] 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.

[16] 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.

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 B.6, and figure B.7 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 caste (figure B.8).

Fact 5: LCs have a higher retention rate at the firm in status quo, owing primarily to worse outisde options. At baseline, there is a high rate of churn among the workers at the firm: 11% on average per month. This is in line with the high rate of worker exits in manufacturing in developing countries in particular. Prior meta-analysis work shows that worker turnover leads to substantial declines in productivity (Hancock et al., 2013). In my setting, 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 B.10). 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).[17] 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.

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

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, 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 A.9.

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.[18] This design is represented in figure A.2.

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 A.1. 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

[18]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 A.3.

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 A.4). 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 to vary the extent of coordination across teams.[19] 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

[19] A pictorial representation of the tasks is in figures A.6 and A.7.

order to arrive at a revealed measure of their cohesion, modeled on Breza et al. (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.

4 Empirical 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

The firm's referral allocation problem is dynamic: supervisors decide each period how to allocate referral opportunities among incumbent workers, trading off possible short-run cohesion costs from changing team composition against longer-run gains from retention and stability. This section provides a unified overview that (i) introduces a common output–profit framework used throughout, (ii) situates a one-period static benchmark, (iii) summarizes the dynamic Bayesian learning channel that the experiment activates, and (iv) links these pieces to estimate the firm's perceived profits over time, in order to help rationalize the initial pre-experimental status quo. Full details and estimation are in Appendix G.

5.1 A common output–profit framework

Let a team in period t have state

$$S_t \equiv \left(n_{\ell,t}, n_{u,t}, \bar{q}_{\ell,t}, \bar{q}_{u,t}, \nu_{\kappa,t}, \sigma^2_{\kappa,t}\right),$$

where $n_{g,t}$ counts workers of group $g \in \{\ell, u\}$ (lower caste and upper caste), $\bar{q}_{g,t}$ is the average network quality (a proxy for the workers' outside option), and $(\nu_{\kappa,t}, \sigma^2_{\kappa,t})$ are belief parameters about cohesion costs. Let $L_t = n_{\ell,t} + n_{u,t}$, the LC share $s_t \equiv n_{\ell,t}/L_t$, and the referral share to LCs be $\rho_t \in [0, 1]$.

Production follows a Cobb–Douglas technology augmented by a cohesion term (Bartelsman et al., 2013; Bloom et al., 2016; Hamilton et al., 2003b):

$$Y_t = A \cdot C(S_t, \rho_t)\, L_t^{\alpha} K_t^{\beta}, \qquad C(S_t, \rho_t) \equiv 1 - \kappa_t \, |\rho_t - s_t|^{\eta}, \quad \eta > 0, \tag{6}$$

where κ_t is the (possibly time-varying) cohesion sensitivity and $C(\cdot) \in (0, 1]$ scales team efficiency as referral allocations deviate from the team's composition. Turnover creates training and reorganization costs; let Q_t denote quits and $\mu > 0$ convert quits into monetary costs. I write the reduced-form per-period profit as

$$\pi_t = Y_t - wL_t - \mu Q_t - (\text{implementation frictions}). \tag{7}$$

This decomposition nests the static benchmark and the dynamic model below; Appendix G gives the full specification used for estimation and the mapping to moments.

5.2 Static benchmark (one period)

As a simplified reference point, consider a static one-period version of the firm's problem. This static benchmark serves primarily as a simplified reference to highlight the immediate trade-offs between cohesion and retention, absent the complicating dynamics of belief updating. A firm of fixed size L allocates a share $\rho \in [0, 1]$ of referral opportunities to LC incumbents; the remainder $1 - \rho$ goes to UC incumbents. Current diversity, exogenously determined, is $S \in [0, 1]$, equivalent to the pre-referral LC share of the team.

Production. Output obeys

$$Y \;=\; A\, C(S, \rho)\, L^{\alpha}, \qquad C(S, \rho) \;=\; 1 - \kappa \, |\rho - S|^{\eta}, \quad \kappa > 0, \; \eta > 1,$$

so cohesion (and hence productivity) is maximised when referrals mirror the team's exogenously determined caste composition.

Training cost. Let the inverse turnover (retention) rates be γ_l and γ_u for LCs and UCs, and define the retention advantage $\delta \equiv \gamma_l - \gamma_u > 0$. Expected quits therefore cost the firm:

$$T(\rho) \;=\; \tau\Big(\rho \frac{1}{\gamma_l} + (1 - \rho) \frac{1}{\gamma_u}\Big),$$

with $\tau > 0$ converting exits into monetary terms.

Firm objective. Taking the uniform wage w as given, profit is

$$\pi(\rho; S, \delta) = AL^{\alpha} C(S, \rho) - T(\rho) - wL,$$

and the first-order condition balances the marginal cohesion loss against the marginal training gain.

Closed-form optimum. Solving the first order condition yields an interior solution when-

ever $\delta \neq 0$:

$$\rho^* \;=\; S \;+\; \frac{\delta}{|\delta|}\cdot,\left[\frac{|\,\tau\delta\,|}{AL^{\alpha}\kappa\eta\,\gamma_u(\gamma_u+\delta)}\right]^{\frac{1}{\eta-1}} \quad (\text{if } |\rho^*-S|<1).$$

When $\delta = 0$ the optimum collapses to the boundary $\rho^* = S$.

There are several notable insights from this static model. First, the optimal referral share moves in step with the team's exogenously determined LC share. It is optimal for the firm to allocate more referrals to LCs in teams that have more LCs to begin with. Second, the share of referrals to LCs grows with a higher retention advantage to LCs (i.e. larger δ). As LCs' effect on the team's expected turnover increases owing to a greater difference between their and UCs' average retention rates, the team finds it more profitable to allocate more referrals to LCs. Third, the previous two relationships are mediated by the effect of LC hires on the team's cohesion. When there is a positive retention advantage from hiring LCs (i.e. $\delta > 0$), it becomes optimal for the firm to hire fewer LCs when the team's cohesion is more sensitive to diversity (i.e. larger κ). Fourth, the elasticity of the deviation from exogenous diversity is non-linear. The optimal share of LC referrals increases for small values of η, and then falls when cohesion is sharply misaligned.[20]

This static benchmark indicates that even large training-cost savings (δ) seldom justify allocating *all* referrals to LCs; the profit-maximising firm strikes a balance where the marginal cohesion cost equals the marginal training benefit. These insights and comparative statics are represented in Figures F.1 - F.4 of Appendix F.

5.3 Dynamic Bayesian learning

In the full dynamic Bayesian learning model, supervisors begin with Gaussian priors over the cohesion sensitivity,

$$\kappa \sim \mathcal{N}(\nu_{\kappa,0},\, \sigma^2_{\kappa,0}),$$

and learn from realized output as referral allocations change team composition. In each period with $N_t > 0$ LC referrals, the supervisor observes a per-referral signal

$$s_t \;\equiv\; \frac{Y_t - \bar{Y}_0}{N_t},$$

where $\bar{Y}_0$ is baseline output. Under the production mapping in (6), $s_t \mid \kappa \sim \mathcal{N}(\kappa,\, \sigma^2_\varepsilon/N_t)$, yielding conjugate Normal–Normal updates for $(\nu_{\kappa,t},\, \sigma^2_{\kappa,t})$ once a minimal stock of LC hires is accumulated. Mandatory LC referrals in treatment ensure that threshold is crossed quickly; control teams rarely accumulate enough exposure to update beliefs materially. Figure 3 documents the resulting divergence in inferred cohesion costs across arms.

[20] The proofs for these comparative statics are in Appendix F.

5.4 Perceived profits and costly experimentation

To connect learning and worker flows to bottom-line firm performance, I simulate *perceived* weekly profits using the same panel objects as the reduced-form analysis. The simulator takes monthly entry flows from the data, constructs tenure-specific exit hazards for LCs and UCs, allows for a short delay before entrants become productive, and embeds a parsimonious belief path ν_t that drifts from a pessimistic prior κ_{prior} toward the true κ_{true} once sufficient evidence accumulates. Weekly output is then

$$y_t \;=\; A\left(\omega L_t^{\ell} + (1-\omega)L_t^{u}\right)^{\alpha} \cdot \phi(\nu_t) + \varepsilon_t, \quad \phi(\nu_t) = \frac{1}{1+\phi_{\text{belief}}\nu_t} \in (0,1],$$

and profits are

$$\pi_t = y_t - \mu C_t - f - \mathbf{1}\{t \geq 1\}\, c_0 e^{-t/\tau_{\text{impl}}} - \lambda\,|\rho_t - s_t|, \tag{8}$$

where C_t is churn, f a fixed cost, c_0 and τ_{impl} capture a fading implementation friction, and $\lambda|\rho_t - s_t|$ penalizes short-run reorganization as the referral rule departs from the status quo. I estimate parameters by simulated method of moments (SMM) against three sets of moments: (i) event-time differences in standardized weekly output, (ii) monthly entry and exit (i.e. churn), and (iii) monthly referral allocations and resulting LC shares in team composition. Figure 4 reports the simulated path for perceived firm profits under the treatment policy, alongwith 95% confidence intervals arrived at through Monte Carlo simulations of the estimation process. The treatment path exhibits a short initial dip due to (a) pessimistic beliefs, (b) implementation frictions, and (c) churn while entrants are delayed on the bench. The profits trajectory then recovers as beliefs improve and LC cohorts mature into productive employment. The mean crossing time (profits > baseline) is about 6 weeks, with a 10–90 percentile interval of 5–14 weeks.

Figures 3 and 4 summarize the mechanisms that rationalize both the pre-experiment status quo and the persistence we observe after the policy is lifted. Treatment acts as *directed exploration*: by mandating LC referrals, it generates the informative draws that move supervisors' posteriors about cohesion costs, pushing beliefs down toward the true (lower) cost in Figure 3. Combined with the retention advantage of LC hires and the resulting stabilization of teams, this learning translates into perceived profits that dip briefly and then turn positive within the first two months in Figure 4. Seen through the lens of the literature on multi-armed bandit problems, the supervisor each period chooses between "arms" that differ in both expected payoff and uncertainty – allocating referrals to UCs (an arm with familiar, myopically safer returns under pessimistic priors) versus to LCs (an arm whose payoff is initially uncertain because cohesion costs are unknown and expected to be high). In the status quo equilibrium, a pessimistic prior, implementation frictions, and churn costs tilt the firm against the LC arm, leading to under-exploration. The intervention forces the LC arm to be pulled often enough to clear a stock threshold for learning, shrink posterior variance, and lower the posterior mean cohesion cost. Once these conditions are met, the perceived value of the LC arm rises above that of the UC arm, yielding both the profit path in 4 and the post-experiment persistence in LC allocations.

Three testable implications follow from this logic: (i) larger post-experiment LC referral shares where there was more LC exposure (i.e. stronger learning in treatment teams), (ii) a short transitional profit dip driven by implementation frictions, benching and churn, followed by recovery as beliefs improve and LC cohorts mature, and (iii) stronger gains in tasks with high coordination intensity where stability is most valuable. The formal state-space, belief updates, and the profit mapping used to estimate these dynamics are in Appendix G.

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.[21] 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%.[22]. Importantly, these changes in the composition are not associated with 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

[21] 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.

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

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.[23] Meanwhile, in control teams, LC referrals drop to rates below the share of LC workers in these teams.[24]

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). This result is economically meaningful in this context. Adhvaryu et al. (2023), in the context of an Indian manufacturing firm, show that only about 50% of the firm's workers are retained through the period of a year. Adhvaryu et al. (2022) show, also in the Indian context, that quit rates are around 10% in their context within a quarter. The declines in churn I see in my setting thus represent large savings to the firm. In addition, the treatment effects I find are in line with benchmarks in the personnel economics literature. Hoffman and Tadelis (2021), in the context of a high-tech US firm, find that increasing a manager's people management skills from the 10th to 90th percentile reduces churn by about 60%.

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 B.9).

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

[23] There is growing evidence that firms (or teams within firms, in my setting) must experience a technology before they adopt it. Garg et al. (2025) show that Bangladeshi garment factories only show a willingness to purchase air purifiers after trial periods that let managers see productivity gains first hand; information, credit, and warranty treatments alone had no effect. A similar experiential barrier appears in Atkin et al. (2017), where soccer-ball manufacturers in Pakistan rarely adopt a cost-saving cutting technology even when it is handed to them for free. In the Bloom et al. (2012), Indian textile plants took up new practices only after consultants worked on-site, again pointing to hands-on learning as the driver of adoption.

[24] 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, this masks considerable heterogeneity with respect to the incidence of exposure to LC referrals during the intervention. I show that there are qualitatively more LC referrals in the post-experiment period among control teams that randomly gave more referrals to LCs in the experimental period, though this difference is not always powered owing to a small sample size (figure B.15).

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 already apparent. I show descriptive comparisons between upper and lower caste incumbent workers' referral candidates on four key dimensions in Figure B.9. 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.

Employers need only a modest degree of monopsony power to align hiring and training decisions with the strength of workers' outside options. Dynamic search models with counter-offers show that, once a firm can infer the arrival rate of competing offers, it invests only in employees whose quit risk is low (Postel-Vinay and Robin, 2006). When a statutory minimum wage pins pay to a floor – as in my setting – wage markdowns are off the table, yet some rents can be captured through ex ante screening: firms may seek to admit applicants whose observable traits signal weak outside options and, hence, longer expected tenure. This mechanism echoes the logic in Acemoglu and Pischke (1999), who argue that wage compression pushes employers to recoup training costs not through lower pay but through lower turnover.

Outside options, reservation wages, and quit propensities are usually hidden from employers at the point of hire. In my context, however, caste provides a visible proxy: it is strongly correlated with outside opportunities and immediately observable during recruitment. The experiment shows that supervisors could learn to exploit this signal, selecting lower-caste referrals to secure a retention advantage. This learning is evidenced in the continued allocations to LCs by treatment supervisors after the conclusion of the experiment, shown in figure 2 and discussed in the previous section.

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 0.09 σ increase in output over the duration of the experiment, translating to a 4.6 percentage point improvement in performance. High-coordination tasks showed a stronger effect of 0.10 σ compared to 0.03 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 B.14, I show that improvments in team-level output intensified with more stability in teams. These results replicate broader findings from the literature on referrals across contexts, where referred workers have been shown to drive more productivity and profits for their firms than non-referred workers, driven primarily by lower turnover among referred workers (Burks et al., 2015). In this sense, LC referred workers seem to drive treatment effects on the intensive margin of referral allocation similar to the extensive margin effects of referrals as a recruitment tool themselves.

Importantly, in this setting, the value of retention shows up in output more strongly 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 B.14. 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 B.15).

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 A.9, 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.[25] 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 B.13).

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.

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[26]), output at the individual and team level, and the initial self-reported cohesion. The information reported to the respondent is presented in figure A.8. 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

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

[26] 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.

order. The six recruitment policies presented to respondents are listed in the table below, with the text in italics listing the actual content that was communicated to the respondents.

Recruitment mode	Caste label visible	Communication to respondent
LC referral	No	*The new worker is hired through a referral given to [name of LC incumbent]*
LC outsider	No	*The new worker is unknown to existing workers, and his name is [name of new LC worker]*
UC referral	No	*The new worker is hired through a referral given to [name of a UC incumbent]*
UC outsider	No	*The new worker is unknown to existing workers, and his name is [name of new UC worker]*
LC referral	Yes	*The new worker is hired through a referral given to the lower caste worker: [name of LC incumbent]*
LC outsider	Yes	*The new worker, who 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 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 8, 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 A.8), 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 B.21.

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 6. 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 A.5. 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 B.1, 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 B.1, 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 B.2. 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.

Effects on referral candidates

A key feature of the design of the firm experiment is the fact that I elicit an unconstrained list of referral candidates as part of the workers' baseline surveys. This allows me to focus on a cohort of workers who were not at the firm at baseline in the treatment group and in the control group, as in figure A.2). There is a range of noteworthy questions for this population. First, does the treatment affect the aggregate job search activity for the connections of LC incumbent workers in treatment teams differently from the connections of LC incumbent workers in the random seeding teams? Second, does this displace the job arrival rate for the connections of UC incumbent workers in treatment teams relative to those in the random seeding teams? In order to answer these questions, I conducted a telephonic labor force participation and job search activity survey with the full set of referral candidates approximately 10 months after the baseline, and approximately 4 months after the culmination of the firm experiment.

There are two forms of attrition-related issues I faced in this survey. First, I have phone numbers for only 31% of the referral candidates. As part of the baseline elicitation, I

collected a range of indicators for each of the referral candidates including their education, current employment status, and location. These indicators are available for the universe of referral candidates, and are analyzed and discussed in fact 6 and figure B.9. However, only a subset of incumbent workers are able to provide phone numbers for their candidates.[27] Second, conditional on having phone numbers, 61% of the referral candidates responded to the survey. This is consistent with typical response rates in phone surveys in similar contexts (Muralidharan et al., 2021; Brownstone and Srivastava, 2025). I present a comparison of the sample of incumbent workers who do and do not share phone numbers of referral candidates, and a comparison of referral candidates by a combination of the indicators reflecting whether they had a listed phone number and whether they responded to the survey. I do not find systematic differences between these populations.

In table B.18, I present results on referral candidates' job arrival rates. There is a reallocation in the treatment group on the number of job offers referral candidates had heard about in the preceding six months, decreasing among treated upper caste candidates (0.22 jobs, 16%) while increasing for treated lower caste candidates (0.27, 52%). While these effects are individually noisy, the difference in the estimates for each caste is statistically different from zero ($p = 0.04$). I find directionally similar effects on the probability of getting a job offer from a factory in the last six months, which decreases for upper castes (9%) while increasing for lower castes (26%), with the difference between caste groups significantly different from zero ($p = 0.08$).

In these surveys, I also elicited basic information on realized employment and recent spells of unemployment (table 7). There is no significant treatment effect on the probability of salaried employment among upper caste referral candidates, while the probability for lower caste referral candidates increases significantly (50%, $p < 0.1$). Likewise, there is no detectable effect on the duration of the most recent spell of unemployment for upper caste referral candidates, while there is a 29% decline for lower caste referral candidates ($p < 0.1$).

Overall, these long-run follow-up estimates suggest substantial improvements in the job access and welfare outcomes for lower caste referral candidates, indicating an increased ability to sort into salaried work. These gains come without any detectable costs to the displaced upper caste referral candidates. This is consistent with upper caste candidates having a better stock of network quality and being in the urban labor market already, validated in baseline indicators for these candidates elicited through their connections at the firm.

[27] The primary determinant for whether I have a phone number for the referral candidate is whether the incumbent worker had their phone with them during the baseline survey, and could look up the candidates' phone numbers on their phone's contact book. This, in turn, is highly correlated with the workers' task. For some tasks, supervisors allow their workers to bring their phones to the line, which makes it more likely that they will have their phones during the baseline survey.

Robustness

The main findings are robust to a range of additional checks. First, I implement randomization inference procedures to account for the finite number of clusters and verify that the estimated treatment effects remain statistically significant under the sharp null of no effect. Second, I examine the potential influence of social desirability bias in self-reported cohesion measures using an abridged Marlowe-Crowne social desirability scale as detailed in Dhar et al. (2022). Social desiribility does not seem to be driving responses in the endline surveys at the firm, and there is no difference across treatment groups on this specification. Lastly, Across specifications and other tests such as the exclusion of outliers, the main results remain qualitatively unchanged. Detailed robustness tables and figures are reported in Appendix C.

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 workforce 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 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	All	RCT period	All	RCT period	All
Treatment	0.088***	0.108**	0.090**	0.170***	0.030	-0.036
	(0.032)	(0.042)	(0.036)	(0.058)	(0.028)	(0.030)
Control mean	-0.012	-0.021	0.084	0.043	-0.151	-0.114
High vs low p-val					0.045	0.930
R^2	0.82	0.75	0.88	0.79	0.77	0.73
Observations	3432	5940	2080	3600	1352	2340
# teams	132	132	80	80	52	52
# weeks	26	45	26	45	26	45

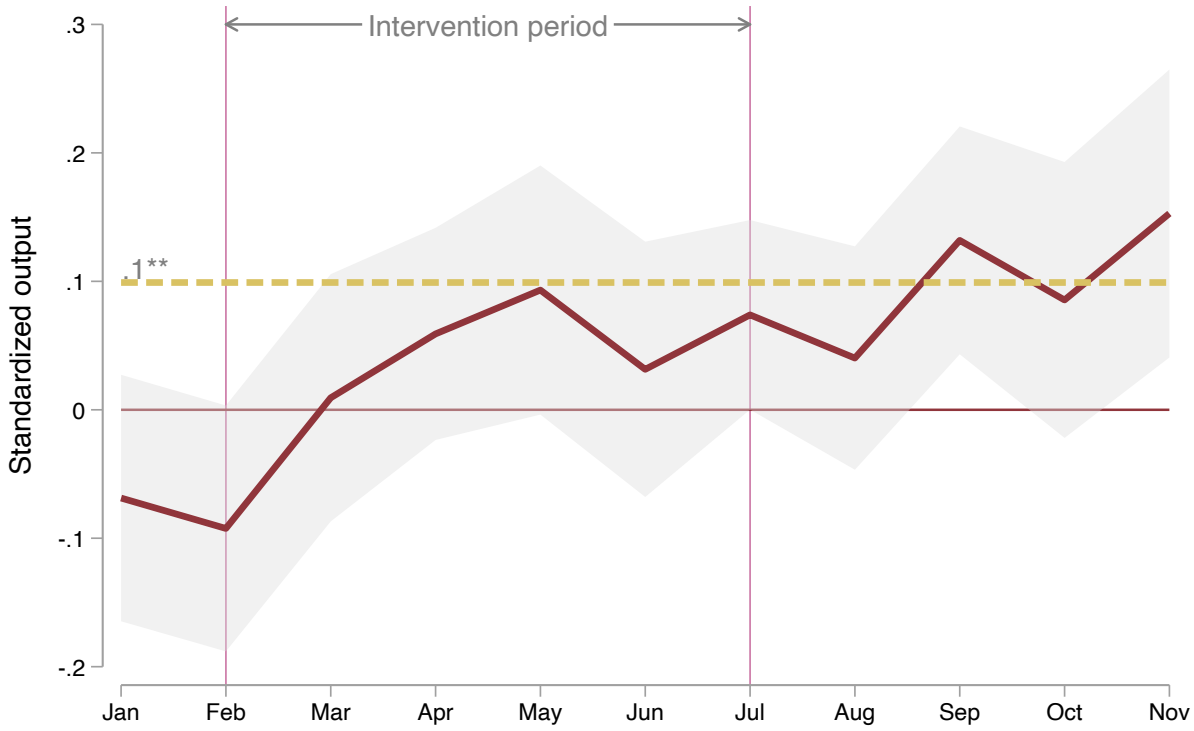
This table uses administrative production data at the team × week level. The output is standardized at the team-type level, using pre-experiment data. Columns 1 and 2 report treatment effects on the standardized output for all teams over the entire period that the data is available for (February - November, 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

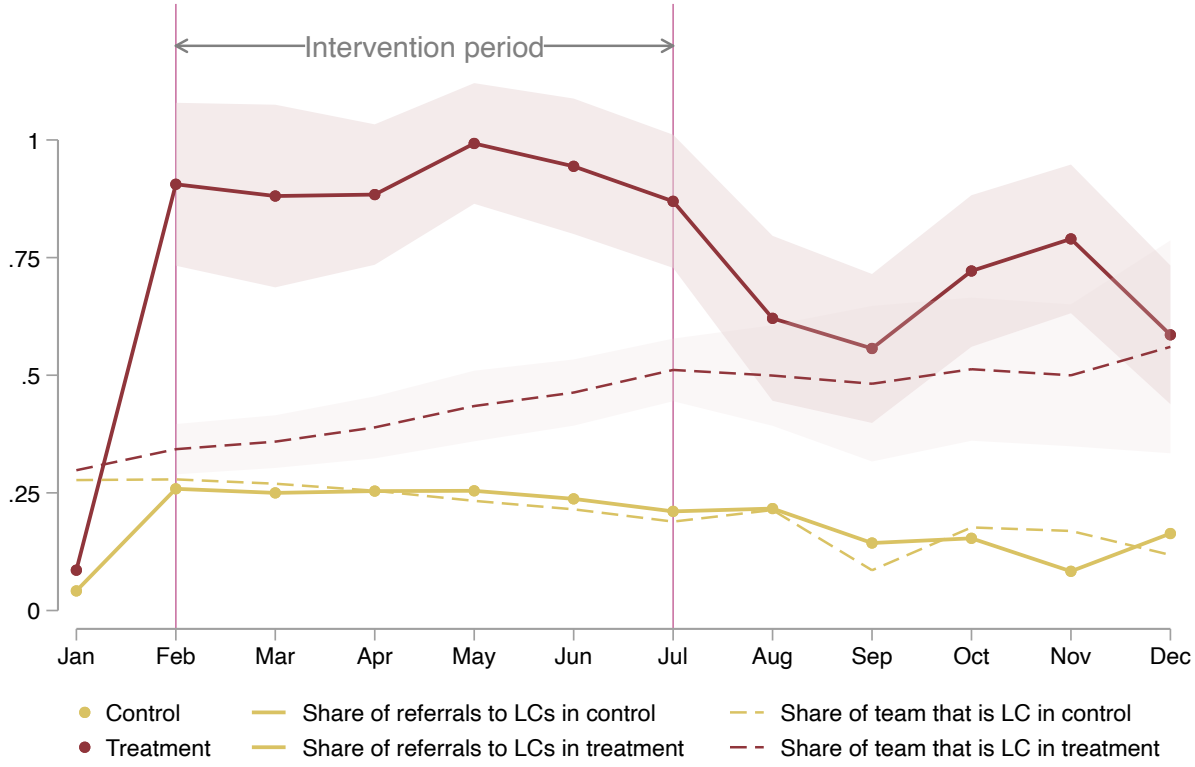
This table uses personnel data at the team × month level to report treatment effects on worker retention outcomes from the experiment at the firm. In columns 1-4, I report the treatment effect on the number of total exits, incumbent exits, exits among new hires, and exits among lower castes respectively. 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 strangers 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 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 stranger 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. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 7: Referral candidates' labor market outcomes

	Pr(in salaried work)		Most recent unemployment spell (in days)	
	(1) Upper caste	(2) Lower caste	(3) Upper caste	(4) Lower caste
Treatment	0.033	0.120*	-1.555	-11.018*
	(0.040)	(0.062)	(3.900)	(6.603)
UC vs. LC p-value		0.238		0.218
Control mean	0.823	0.710	27.440	37.900
R^2	0.00	0.02	0.00	0.01
Observations	418	200	418	185

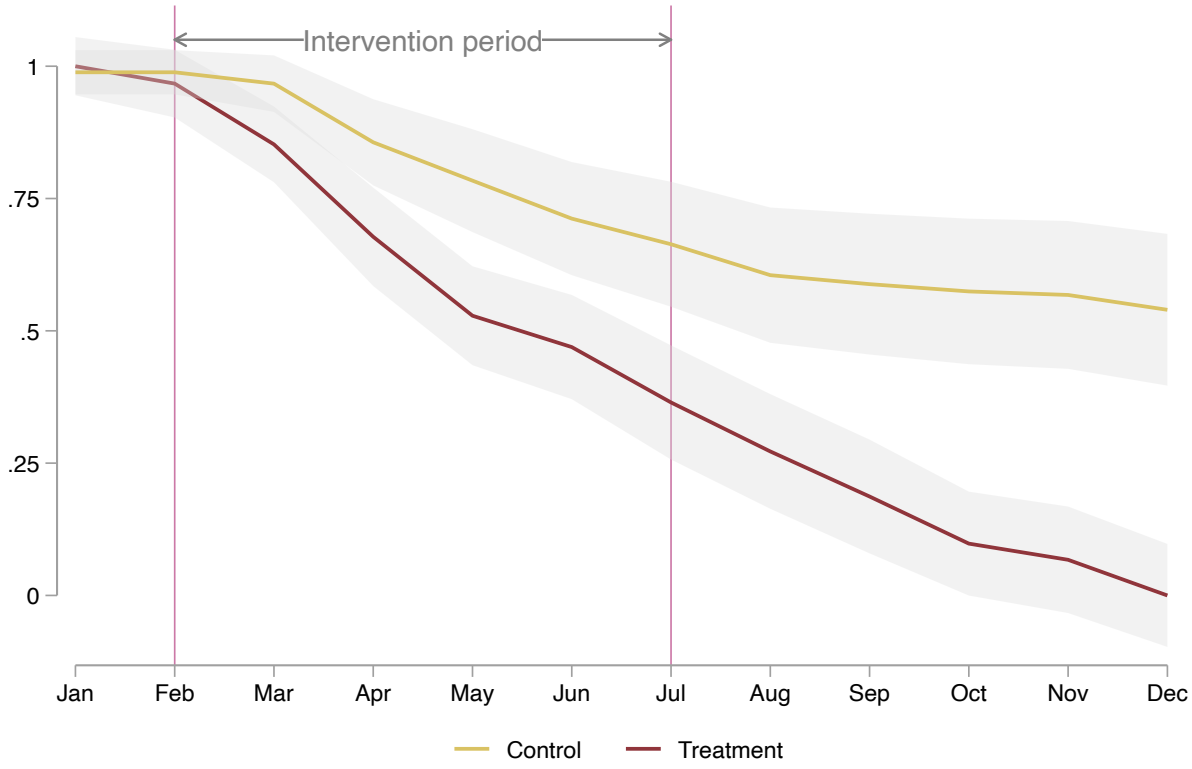
This table reports treatment effects on the labor market activity among referral candidates, i.e. connections of the incumbent workers, as reported in a phone-based follow-up survey. All standard errors are clustered at the team level, and the regressions inlcude fixed effects at the randomization strata level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table 8: Vignettes survey: changes in output

		Change in expected output per worker after recruitment						Obs
		Mean	Pairwise test [row - column]					
Recruitment type ↓→	Caste label ↓→		UC stranger No	UC referral No	LC stranger No	LC referral No	LC stranger Yes	
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
UC stranger	No	0.16	-					93
UC referral	No	0.48	-0.09	-				107
			(0.30)					
LC stranger	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 stranger	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)	

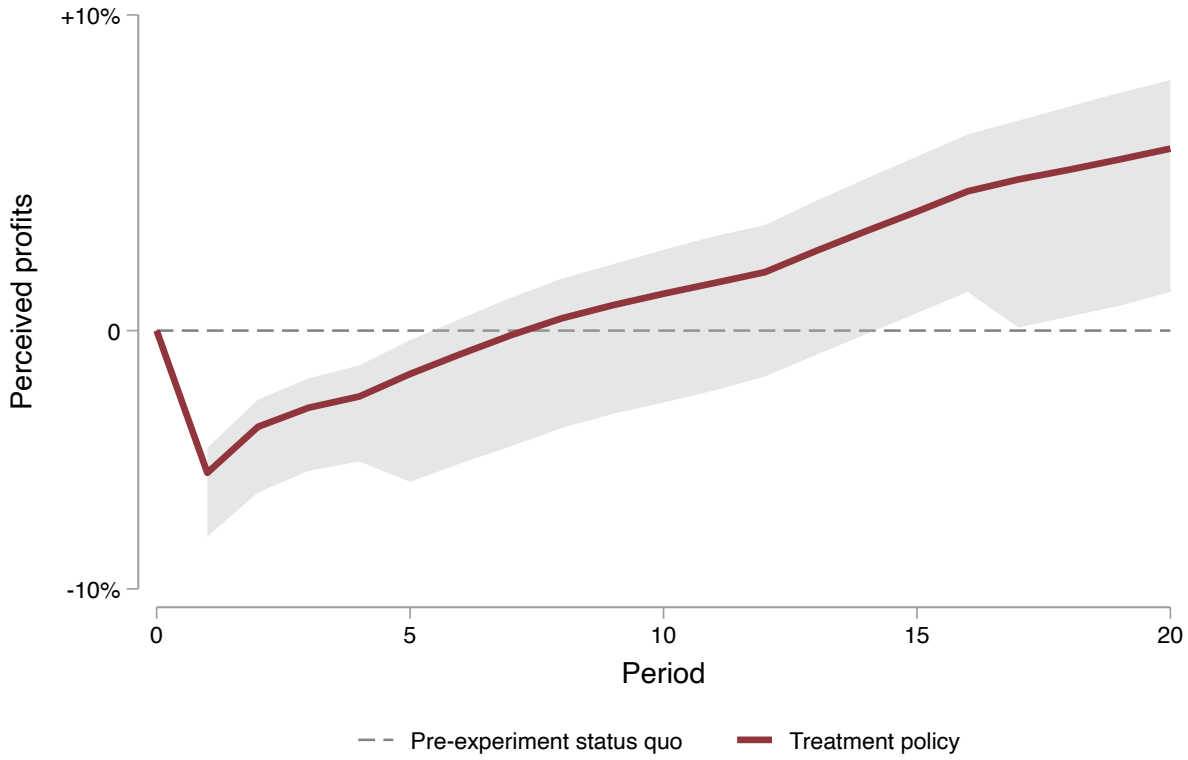
This table reports average reported responses for the change in expected output per worker as a hypothetical team grows from three to four through each reported mode of recruitment. The modes are listed in the labels of rows and columns, with each cell of the matrix in columns 2-6 reporting the comparison 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. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Figure 3: Evolution of beliefs on cohesion costs



This figure plots the trajectory of supervisor beliefs about κ, i.e. the cohesion cost in the Bayesian learning model described in section 5. The bands around treatment and control team trajectories indicate 95% confidence intervals. The initial beliefs are from the baseline survey with supervisors, and the trajectories are estimated from team-level panel data through the process described in section 5.

Figure 4: Evolution of perceived profits



This figure plots the simulated trajectory of perceived profits under the treatment policy relative to the pre-experiment benchmark, normalized to the mean of the control arm in the pre-intervention period. The solid red line shows the treatment path; the shaded region is the 95% Monte Carlo envelope from binomial exit and Gaussian output shocks. The initial dip reflects implementation frictions, pessimistic priors about cohesion, and churn while new hires are on the bench. Recovery occurs as supervisors update beliefs about cohesion costs, LC cohorts mature into productive workers, and team stability improves. Profits cross above baseline within 5–14 weeks (mean $\approx$ 6 weeks).

References

Abraham, Rosa and Surbhi Kesar (2025). "Trajectories of labour market transitions in the Indian economy". In: *World Development* 195, p. 107029. ISSN: 0305-750X. DOI: https://doi.org/10.1016/j.worlddev.2025.107029. URL: https://www.sciencedirect.com/science/article/pii/S0305750X25001147.

Acemoglu, Daron and Jörn-Steffen Pischke (June 1999). "The Structure of Wages and Investment in General Training". In: 107.3, pp. 539–572.

Adhvaryu, Achyuta, Smit Gade, Teresa Molina, and Anant Nyshadham (2022). "Sotto Voce: The Impacts of Technology to Enhance Worker Voice". In: *Working Paper*.

Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham (2023). "Returns to On-the-Job Soft Skills Training". In: *Journal of Political Economy* 131.8, pp. 2165–2208. DOI: 10.1086/724320. eprint: https://doi.org/10.1086/724320. URL: https://doi.org/10.1086/724320.

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, and Sharath Sharma (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.

Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien (1991). "Optimal Learning by Experimentation". In: *The Review of Economic Studies* 58.4, pp. 621–654. ISSN: 00346527, 1467937X. URL: http://www.jstor.org/stable/2297825 (visited on 05/10/2025).

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

Álvarez Pereira, Brais, Shan Aman Rana, and Alexia Delfino (2024). "Team Size and Diversity". In: *Journal of Economic Behavior & Organization* 216, pp. 220–246. DOI: 10.1016/j.jebo.2024.04.004.

Anders, Jenna and Amanda Pallais (2025). "Inefficient Social Disconnection". In: *Working Paper*.

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.

Anisfeld, Ari and Pepi Pandiloski (2024). "Language, Identity, and Ethnic Cohesion in Schools". In: *Working Paper*.

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.

Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit K. Khandelwal, and Eric Verhoogen (2017). "Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan". In: *The Quarterly Journal of Economics* 132.3, pp. 1101–1164. DOI: 10.1093/qje/qjx010.

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

— (2011). "Field Experiments with Firms". In: *Journal of Economic Perspectives* 25.3, pp. 63–82. DOI: 10.1257/jep.25.3.63.

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.

Banerjee, Abhijit and Andrew Newman (1993). "Occupational Choice and the Process of Development". In: *Journal of Political Economy* 101.2, pp. 274–298. DOI: 10.1086/261876.

Banerjee, Biswajit and J. B. Knight (1985). "Caste Discrimination in the Indian Urban Labour Market". In: *Journal of Development Economics* 17.3, pp. 277–307. DOI: 10.1016/0304-3878(85)90094-X.

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

Beaman et al. (2018). "Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi". In: *Journal of Labor Economics*.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts (Nov. 2012). "Does Management Matter? Evidence from India *". In: *The Quarterly Journal of Economics* 128.1, pp. 1–51. ISSN: 0033-5533. DOI: 10.1093/qje/qjs044. eprint: https://academic.oup.com/qje/article-pdf/128/1/1/30624754/qjs044.pdf. URL: https://doi.org/10.1093/qje/qjs044.

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*.

Bolton, Patrick and Christopher Harris (1999). "Strategic Experimentation". In: *Econometrica* 67.2, pp. 349–374. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/2999588 (visited on 05/10/2025).

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.

Brownstone, Steven and Kartik Srivastava (2025). "The long shadow of feudalism: Concentration of land and labor market power in India". In: *Working Paper*.

Bryan, Gharad, Dean Karlan, and Jonathan Zinman (2015). "Referrals: Peer Screening and Enforcement in a Consumer Credit Field Experiment". In: *American Economic Journal: Microeconomics* 7.3, pp. 174–204. DOI: 10.1257/mic.20130234.

Burks, Stephen V., Bo Cowgill, Mitchell Hoffman, and Michael Housman (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*.

Chakraborty, Anujit, Arkadev Ghosh, Matt Lowe, and Gareth Nellis (Sept. 2024). *Learning About Outgroups: The Impact of Broad Versus Deep Interactions*. 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, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez (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.

Dal Bó, Pedro, Andrew Foster, and Louis Putterman (2010). "Institutions and Behavior: Experimental Evidence on the Effects of Democracy". In: *American Economic Review* 100.5, pp. 2205–2229. DOI: 10.1257/aer.100.5.2205.

Dale-Olsen, Harald and Henning Finseraas (2020). "Linguistic Diversity and Workplace Productivity". In: *Labour Economics* 67.

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.

Dhar, Diva, Tarun Jain, and Seema Jayachandran (Mar. 2022). "Reshaping Adolescents' Gender Attitudes: Evidence from a School-Based Experiment in India". In: *American Economic Review* 112.3, pp. 899–927. DOI: 10.1257/aer.20201112. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20201112.

Donovan, Kevin, Wenjie J. Lu, and Todd Schoellman (May 2023). "Labor Market Dynamics and Development". In: *The Quarterly Journal of Economics* 138.4, pp. 2287–2325. ISSN:

0033-5533. DOI: 10.1093/qje/qjad019. eprint: https://academic.oup.com/qje/article-pdf/138/4/2287/51765371/qjad019.pdf. URL: https://doi.org/10.1093/qje/qjad019.

Dustmann, Christian, Alberecht Glitz, Uta Schonberg, and Herbert Brucker (2016). "Referral-based Job Search Networks". In: *The Review of Economic Studies* 83.

Exley, Christine L. and Judd B. Kessler (2024). "Motivated Errors". In: *American Economic Review* 114.4, pp. 961–987. DOI: 10.1257/aer.20191849.

Fadlon, Itzik, Briana Sullivan, and Vedant Vohra (2025). "The Racial Penalty in Job Ladder Transitions".

Fershtman, Chaim and Uri Gneezy (2001). "Discrimination in a Segmented Society: An Experimental Approach". In: *Quarterly Journal of Economics* 116, pp. 351–377.

Finseraas, Henning, Torbjørn Hanson, Åshild A. Johnsen, Andreas Kotsadam, and Gaute Torsvik (2019). "Trust, Ethnic Diversity, and Personal Contact: A Field Experiment". In: *Journal of Public Economics* 173, pp. 72–84. DOI: 10.1016/j.jpubeco.2019.02.002.

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.

Friebel, Guido, Matthias Heinz, Mitchell Hoffman, and Nick Zubanov (2023). "What Do Employee Referral Programs Do? Measuring the Direct and Overall Effects of a Management Practice". In: *Journal of Political Economy* 131.3, pp. 633–686. DOI: 10.1086/721735.

Garg, Teevrat, Maulik Jagnani, and Nancy Lozano-Gracia (2025). "High Returns, Low Adoption: Air Purifiers in Bangladeshi Firms". In: *Working Paper*.

Ghosh, Arkadev (Dec. 2023). "Religious Divisions and Production Technology: Experimental Evidence from India". In: *Working Paper*.

Gittins, John C. and DM Jones (1979). "A dynamic allocation index for the discounted multiarmed bandit problem". In: *Biometrika* 66, pp. 561–565. URL: https://api.semanticscholar.org/CorpusID:119508462.

Granovetter, Mark S (1973). "The strength of weak ties". In: *American journal of sociology* 78.6, pp. 1360–1380.

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.

Hancock, Julie I., David G. Allen, Frank A. Bosco, Karen R. McDaniel, and Charles A. Pierce (2013). "Meta-Analytic Review of Employee Turnover as a Predictor of Firm Performance". In: *Journal of Management* 39.3, pp. 573–603. DOI: 10.1177/0149206311424943.

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.

Hoffman, Mitchell and Christopher T. Stanton (2024). *People, Practices, and Productivity: A Review of New Advances in Personnel Economics*. NBER Working Paper 32849. National Bureau of Economic Research. DOI: 10.3386/w32849.

Hoffman, Mitchell and Steven Tadelis (2021). "People Management Skills, Employee Attrition, and Manager Rewards: An Empirical Analysis". In: *Journal of Political Economy* 129.1, pp. 243–285. DOI: 10.1086/711409. eprint: https://doi.org/10.1086/711409. URL: https://doi.org/10.1086/711409.

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.

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.

Jodhka, Surinder S. and Katherine Newman (2007). "In the Name of Globalisation: Meritocracy, Productivity and the Hidden Language of Caste". In: *The Changing Face of Global India*. Oxford University Press, pp. 135–162.

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.

Li, Danielle, Lindsey Raymond, and Peter Bergman (2025). "Hiring as Exploration". In: *Review of Economic Studies* forthcoming.

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.
Manpower Group (2024). *Talent Shortage Survey*. https://go.manpowergroup.com/talent-shortage.
Mas, Alexandre and Enrico Moretti (Mar. 2009). "Peers at Work". In: *American Economic Review* 99.1, pp. 112–45. DOI: 10.1257/aer.99.1.112. URL: https://www.aeaweb.org/articles?id=10.1257/aer.99.1.112.
Mohanty, Aishwarya (Dec. 2024). *How Caste Shapes Migration From Ganjam*. News article. IndiaSpend. URL: https://www.indiaspend.com/replace-with-actual-url (visited on 08/07/2025).
Montgomery, D (1991). "Social Networks and Labor-Market Outcomes: Toward an Economic Analysis". In: *American Economic Review* 81.
Moorjani, Priya, Kumarasamy Thangaraj, Nick Patterson, Mark Lipson, Po-Ru Loh, Periasamy Govindaraj, Bonnie Berger, David Reich, and Lalji Singh (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.
Muralidharan, Karthik, Paul Niehaus, Sandip Sukhtankar, and Jeffrey Weaver (Apr. 2021). "Improving Last-Mile Service Delivery Using Phone-Based Monitoring". In: *American Economic Journal: Applied Economics* 13.2, pp. 52–82.
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.
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.

Pérez Cervera, Laura (2023). "Affirming a Racial Divide? The Political Consequences of Affirmative Action in Brazil". Job-market paper, December 2023. URL: https://lauraperezcervera.weebly.com/research.html.

Postel-Vinay, Fabien and Jean-Marc Robin (2006). "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity". In: *Econometrica* 70.6, pp. 2295–2350. DOI: 10.1111/j.1468-0262.2002.00441.x.

Prat, Andrea (2002). "Should a Team Be Homogeneous?" In: *European Economic Review* 46, pp. 1187–1207.

Roethlisberger, FJ and William Dickinson (1939). *Management and the Worker*. Harvard University Press.

Sapin, Marlène, Dominique Joye, Christof Wolf, Johannes Andersen, Yanjie Bian, Ali Carkoglu, Yang-Chi Fu, Ersin Kalaycioglu, Peter V Marsden, and Tom W Smith (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*.

World Bank (2013). "World Development Report 2013: Jobs". In: *World Bank*.

Xu, Janet (2021). "Prize or Penalty? The Reputational Effects of Diversity Scholarships in the Labor Market". In: *Academy of Management Proceedings* 2021.1, p. 16502. DOI: 10.5465/AMBPP.2021.16502abstract. eprint: https://doi.org/10.5465/AMBPP.2021.16502abstract. URL: https://doi.org/10.5465/AMBPP.2021.16502abstract.

Appendix

A Study implementation details **44**

B Supplementary exhibits **55**
B.1 Original analysis of secondary data 55
B.2 Baseline firm survey 60
B.3 Lab-in-field survey 66
B.4 Treatment effects on firm administrative data 72
B.5 Treatment effects on endline survey at firm 76
B.6 Treatment effects on referral candidate follow-ups 78
B.7 Vignettes survey 82

C Robustness tests **94**
C.1 Randomization inference 94
C.2 Social desirability bias 99

D Attrition and non-response **102**

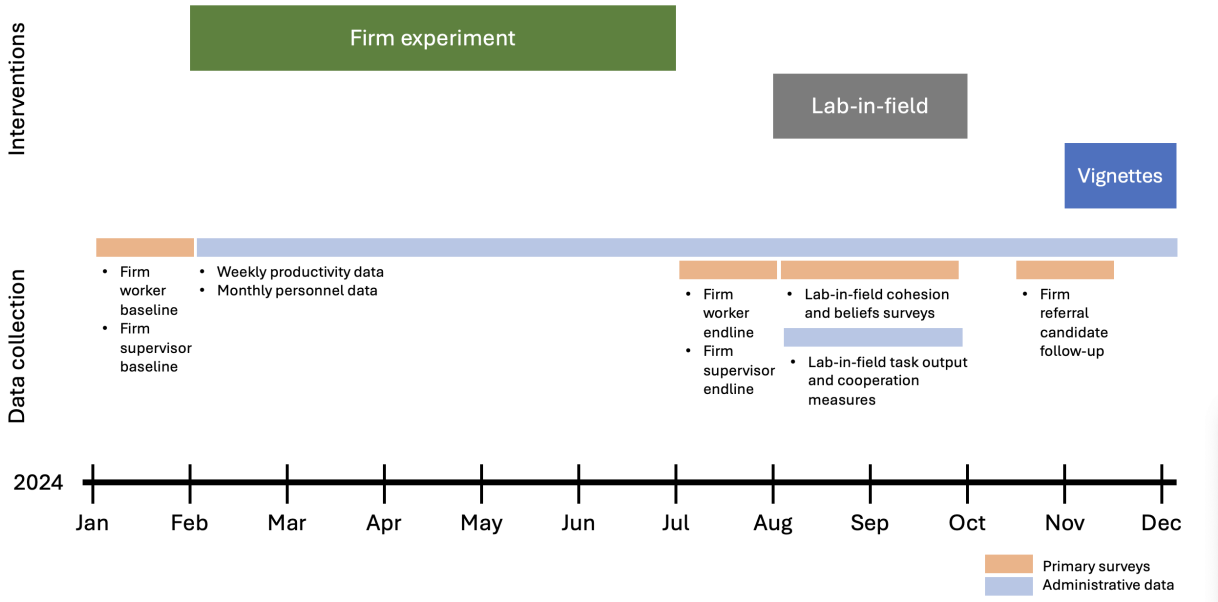
E Deviations from pre-analysis plan **103**

F Static model **105**

G Bayesian learning and perceived profits **114**
G.1 Environment, state and flows 114
G.2 Production and cohesion 114
G.3 Beliefs about cohesion costs 115
G.4 Profits and frictions 115
G.5 Estimation by simulated method of moments 116

A Study implementation details

Figure A.1: 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.

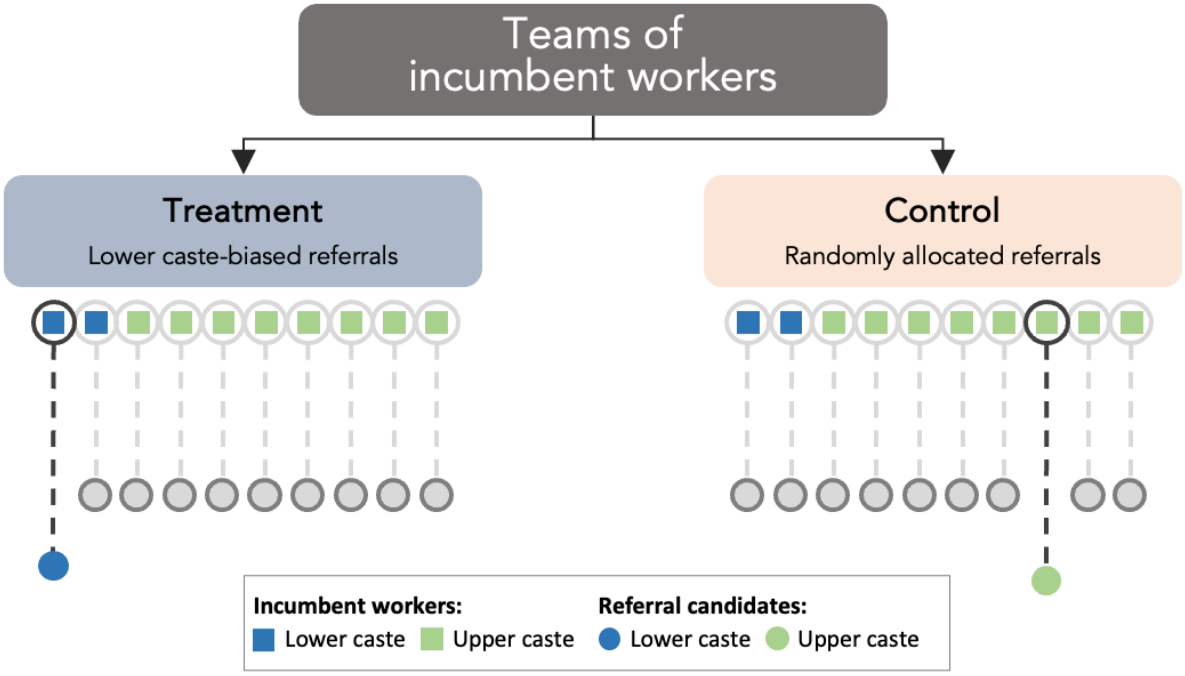
Table A.1: Data sources

Indicators	**Frequency**	**Unit**	**Time**	**Source**
Panel A: Firm experiment				
Output	Weekly	Team	Jan – Dec	Admin data
Entries and exits	Monthly	Individual		
Referral allocations	Monthly	Individual		
Referral candidates elicitation	Once	Individual	Jan	Baseline
Outgroup beliefs	Twice	Individual	Jan & Jul	Baseline & endline
Job satisfaction	Twice	Individual		
Self-reported cohesion	Twice	Individual		
Panel B: Follow-up of referral candidates				
Earnings	Once	Individual	Oct	Referral-candidate endline
Job search activity	Once	Individual		
Panel C: Lab-in-field experiment				
Output	Twice	Team	Aug	Enumerator counts
Symbol matching score	Twice	Individual		
Revealed measures of cohesion	Once	Individual		Endline
Outgroup beliefs	Twice	Individual		
Panel D: Vignettes experiment				
Effect on expected output	Once	Individual	Nov	Vignettes survey
Effect on expected cohesion	Once	Individual		

This table represents the frequency, unit, time period and source of the main indicators used for analysis across the firm experiment, the follow-up survey with referral candidates, the lab-in-field experiment, and the vignettes experiment in panels A, B, C, and D respectively.

Firm experiment

Figure A.2: 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 A.3: 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.

Lab-in-field experiment

Production task. Prior to commencing the production tasks in the lab-in-field workshops, each team member was assigned a random cloth color and instructed to cut strips only from that cloth. To prevent batching effects, participants were required to cut and tie strips sequentially rather than stockpiling them. In high-coordination teams, members were required to follow a predetermined color sequence when tying strips together, necessitating close coordination in both cutting and tying. In low-coordination teams, workers cut and tied strips for their own color independently before combining them across teammates at

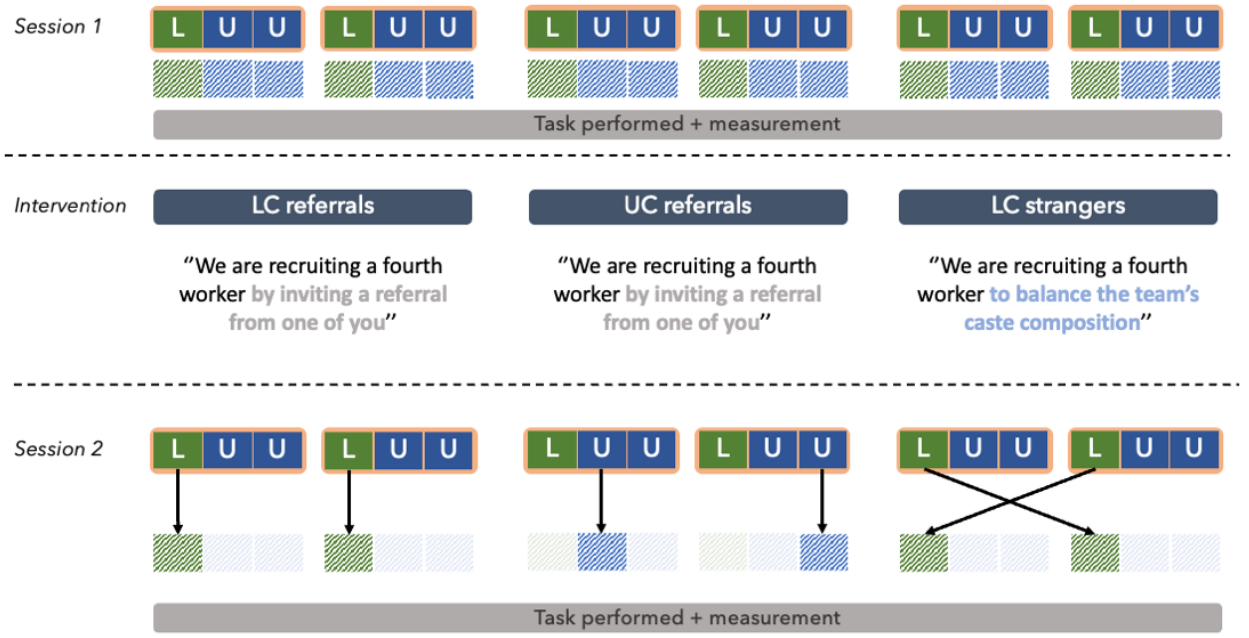
the end.

At the end of each workshop session, two independent enumerators separately measured the length of each finished rope (in the units of the number of strips). Only ropes meeting the pre-specified dimensions (10 cm $\times$ 2 cm per strip) were included in length calculations. The median of the two enumerators' measurements was used as the outcome, with over 90 percent agreement between the two measures. One enumerator also recorded the number of strips of each color to generate disaggregated output measures.

Symbol-matching game. For the symbol-matching game, each worker was randomly paired with a teammate and given five minutes to match as many symbols as possible across their two sheets. Each worker's sheet contained a unique, randomly assigned symbol grid from a set of 15 possible grids. Workers recorded the number of matches they believed to be correct, after which enumerators collected both sheets, noted the IDs of the worker and their partner, and computed the true number of matches. A second round followed with a new randomly assigned teammate.

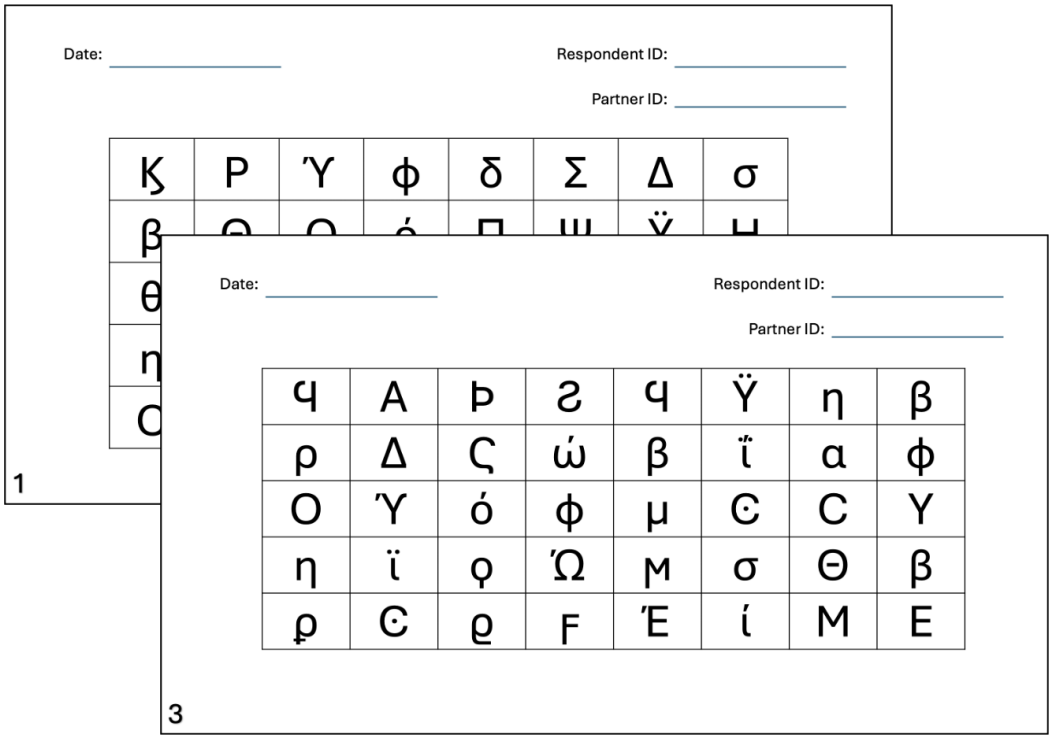
"True" matches were estimated by comparing each combination of the two symbol grids. I then calculated the absolute difference between the true and reported matches, averaged at the worker level, and served as a revealed measure of cohesion. This design ensured independent effort while allowing measurement of coordination quality across varying teammate pairings.

Figure A.4: Lab-in-field experiment design



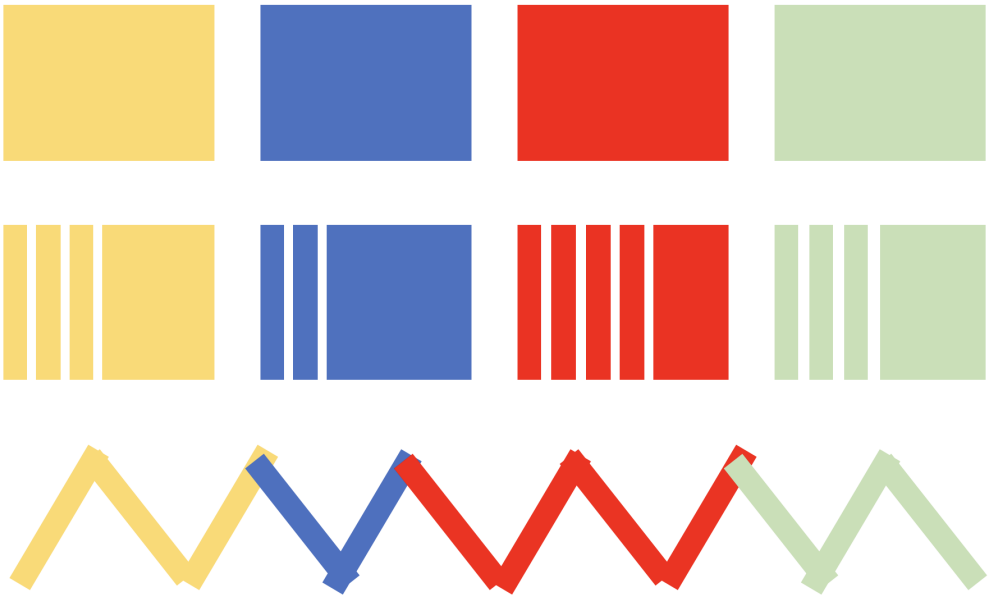
This figure shows the design of the lab-in-field experiment. Each of the workers recruited for this workshop are asked to identify referral candidates. In the first session of the lab-in-field exercise, randomly formed teams with one LC and two UCs each are asked to perform a task. After the production, 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 A.5: Symbol matching game



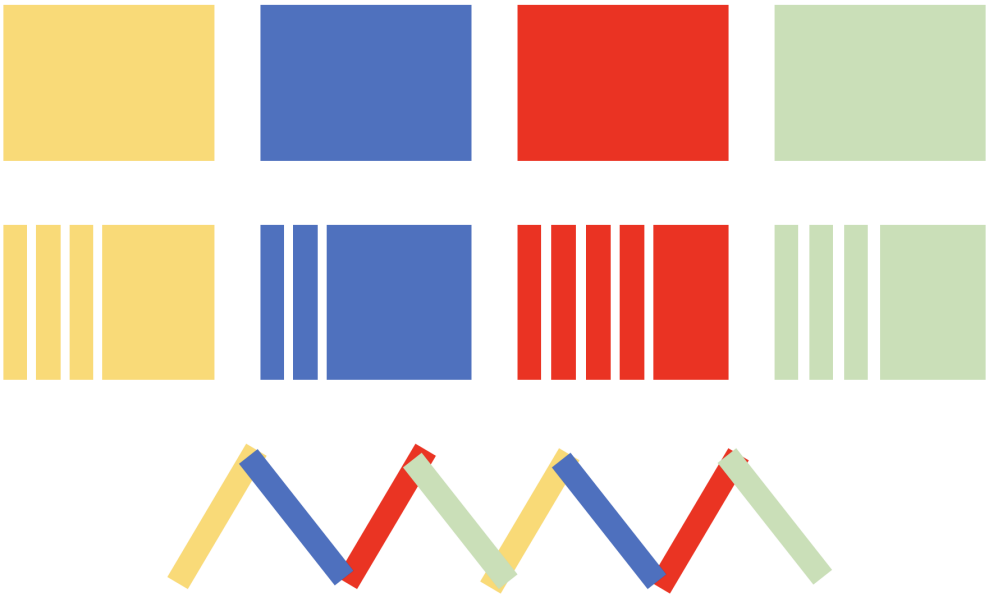
This figure shows an example of the sheets used by the workers in the symbol matching exercise used to get revealed measures of social cohesion.

Figure A.6: 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.

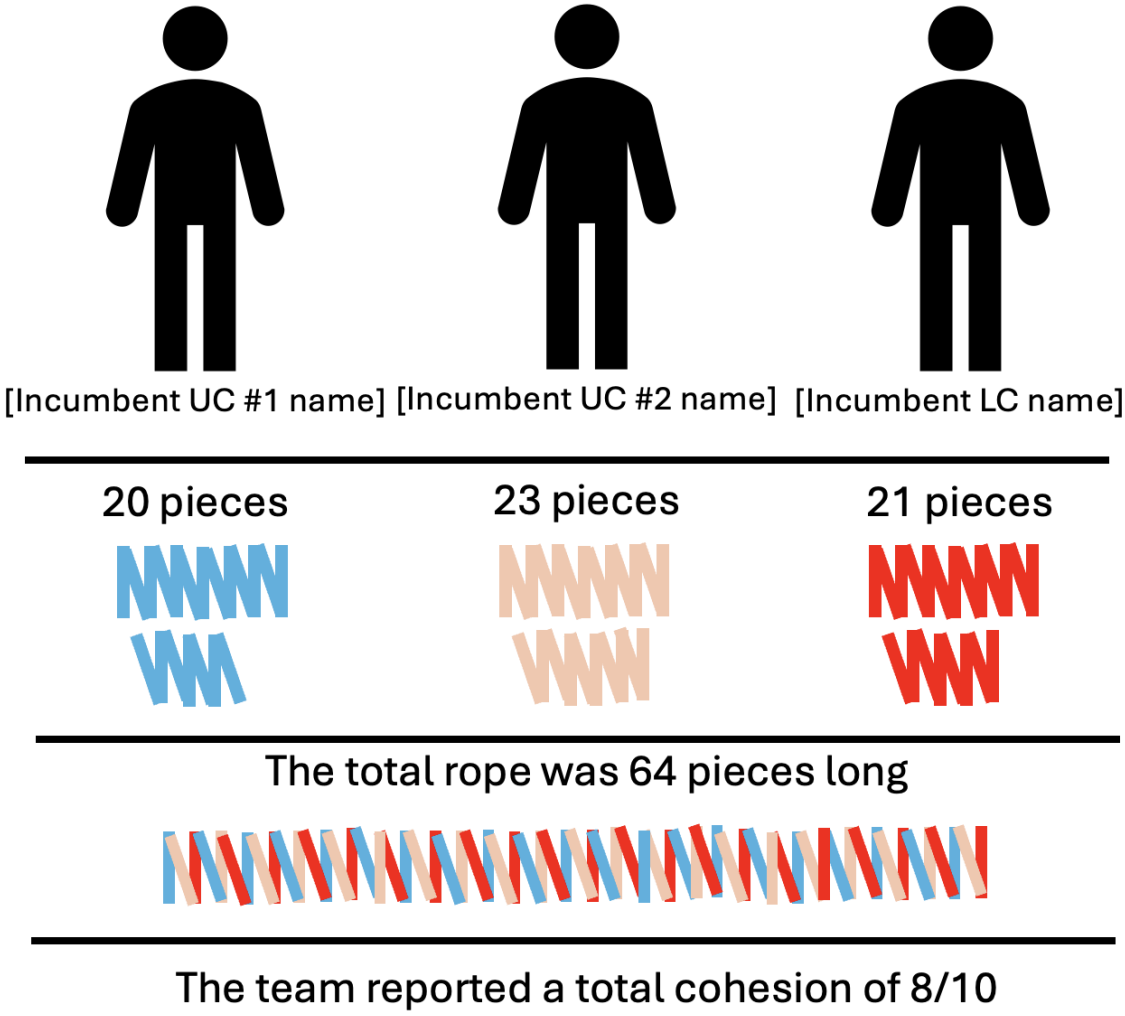
Figure A.7: Illustrative example of the lab-in-field task: high coordination teams



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.

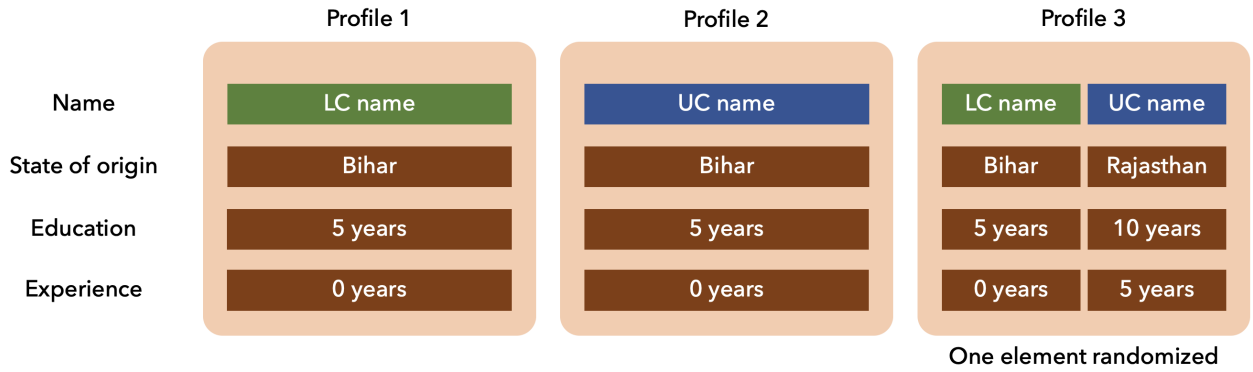
Vignettes experiment

Figure A.8: 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 A.9: Hypothetical profiles for belief elicitation

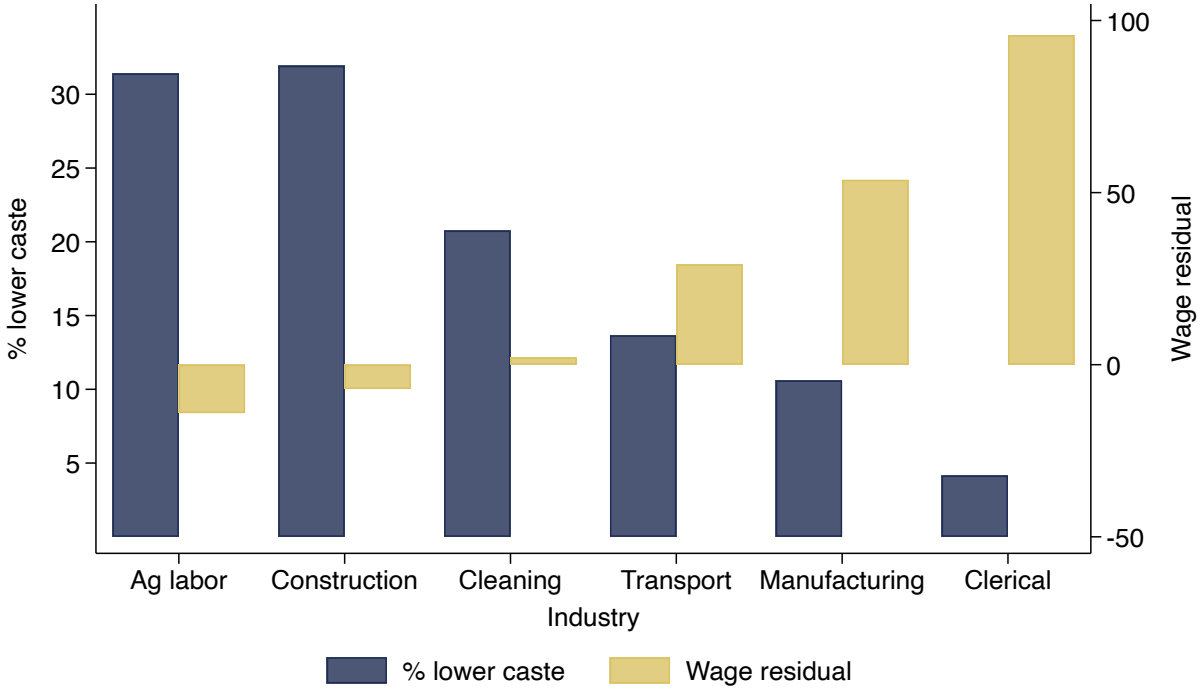


This figure represents the three profiles of hypothetical workers used in the baseline and endline surveys at the firm to elicit worker and supervisor beliefs about out-groups.

B Supplementary exhibits

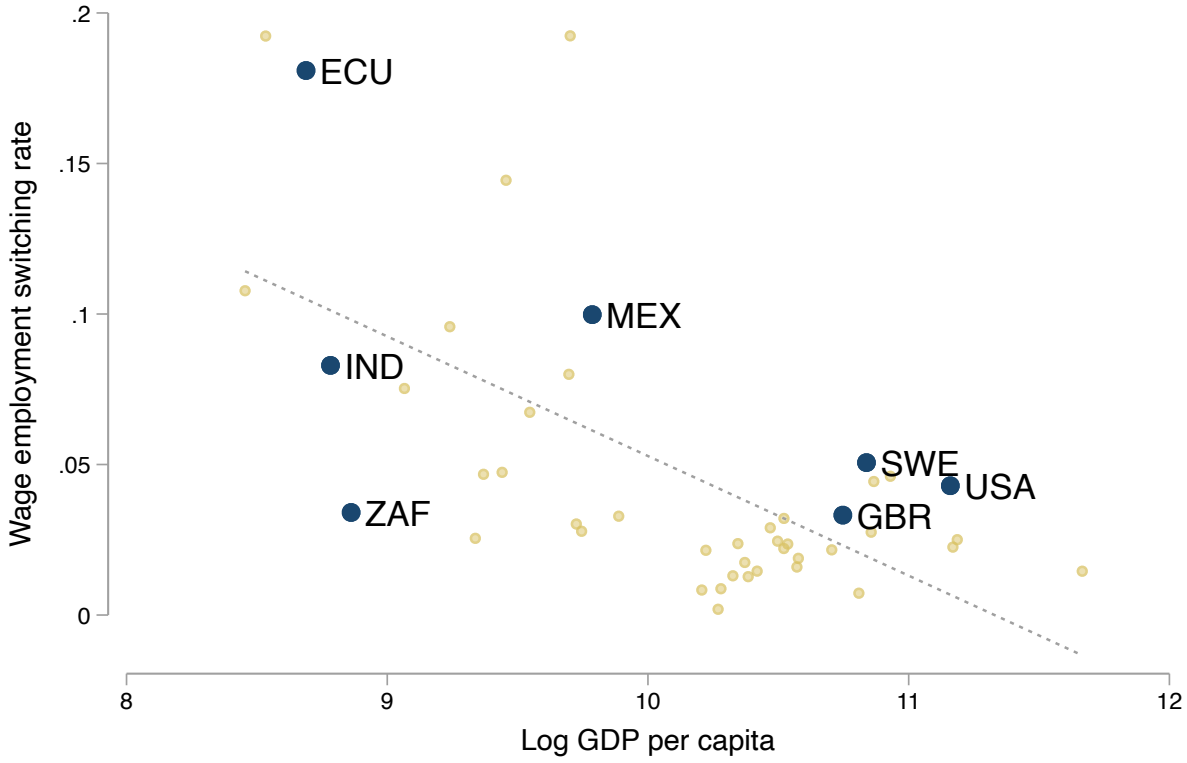
B.1 Original analysis of secondary data

Figure B.1: Descriptive comparison of worker caste composition and industry premia



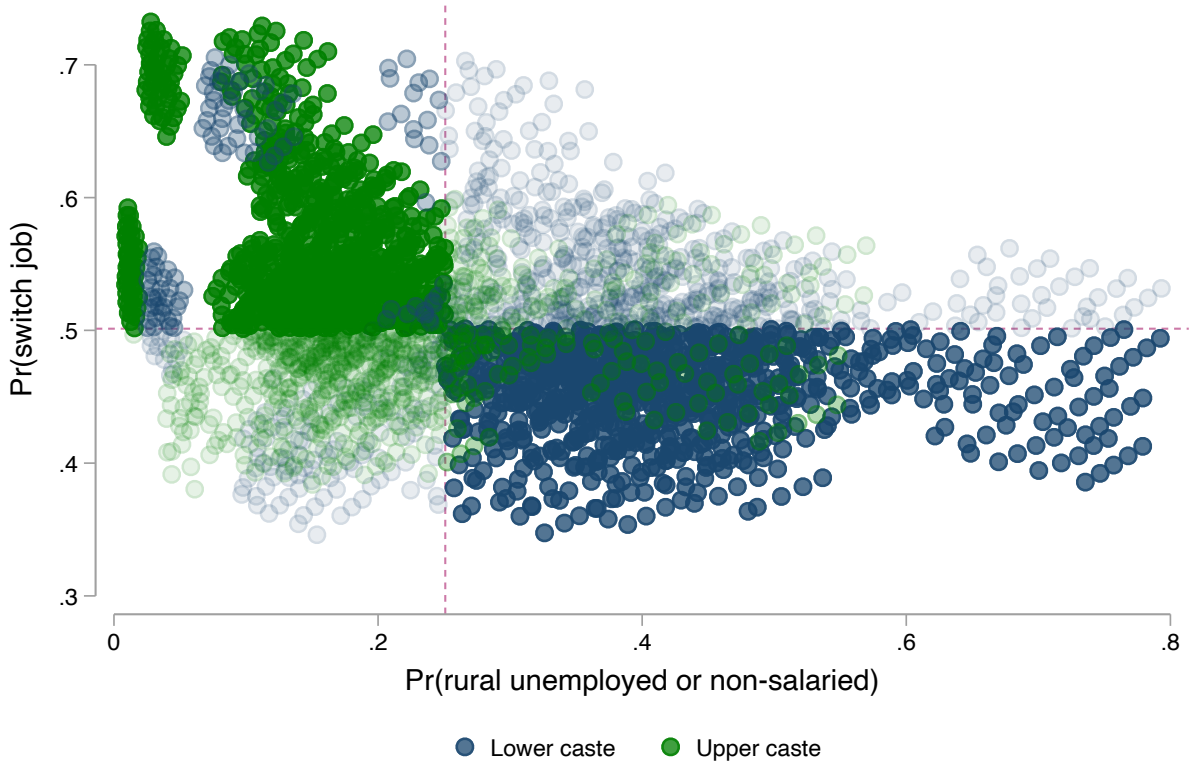
This figure represents original analysis of the nationally representative Indian Human Development Survey 2014-15 data. On the right axis, I plot the wage residual across the lowest key industry-occupation groups reported in the survey, controlling for age, education, gender, and district of birth. On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

Figure B.2: Descriptive comparison of worker caste composition and industry premia



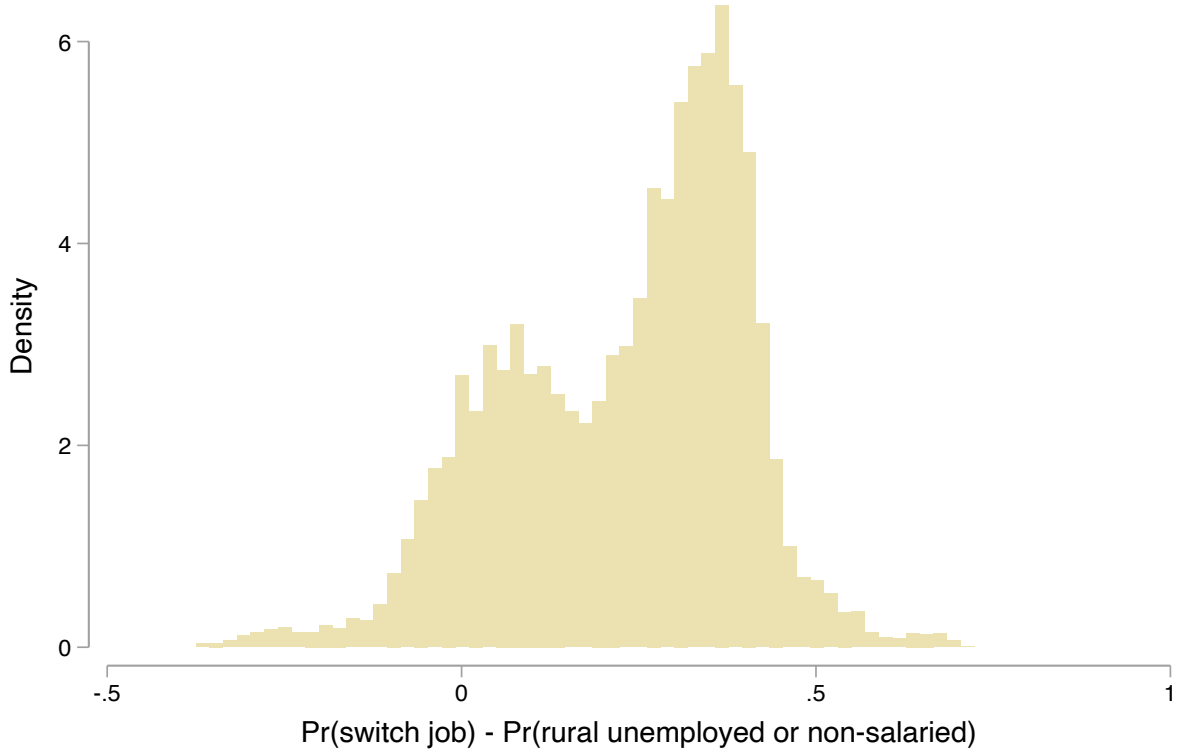
This figure represents original analysis of the nationally representative Indian Human Development Survey 2014-15 data. On the right axis, I plot the wage residual across the lowest key industry-occupation groups reported in the survey, controlling for age, education, gender, and district of birth. On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

Figure B.3: Descriptive comparison of worker caste composition and industry premia



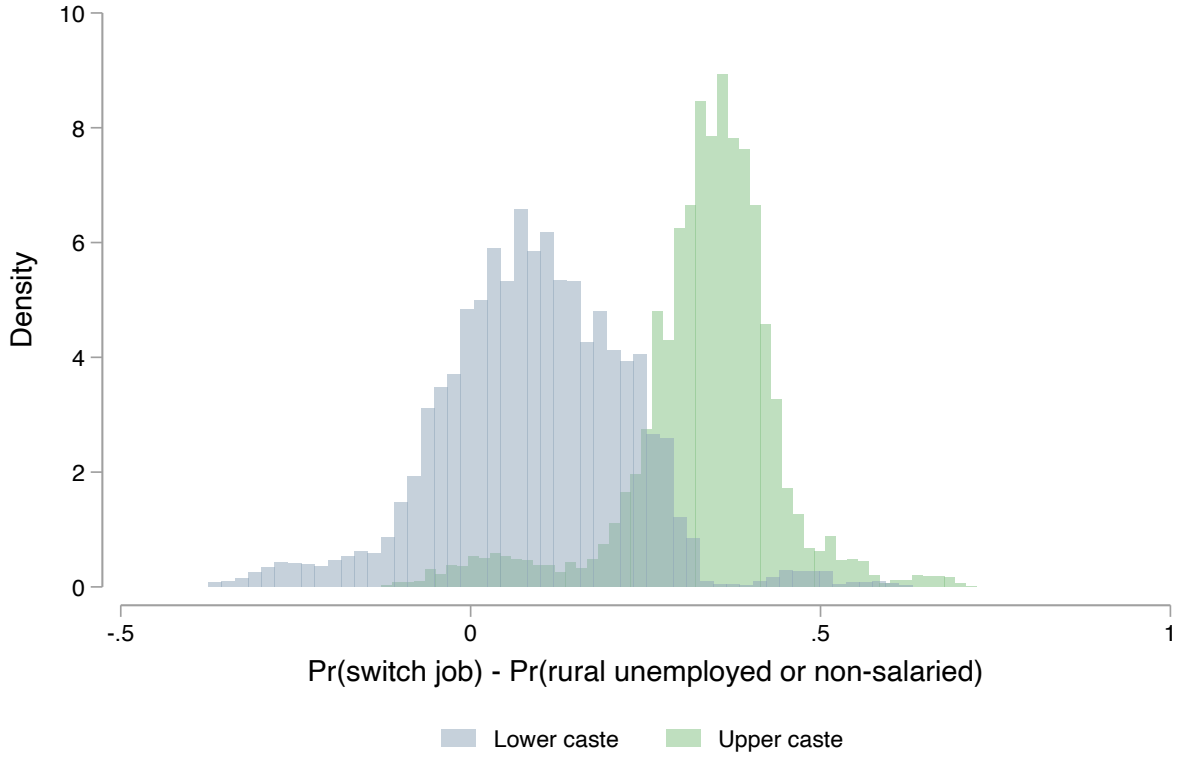
This figure represents original analysis of the nationally representative Indian Human Development Survey 2014-15 data. On the right axis, I plot the wage residual across the lowest key industry-occupation groups reported in the survey, controlling for age, education, gender, and district of birth. On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

Figure B.4: 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.

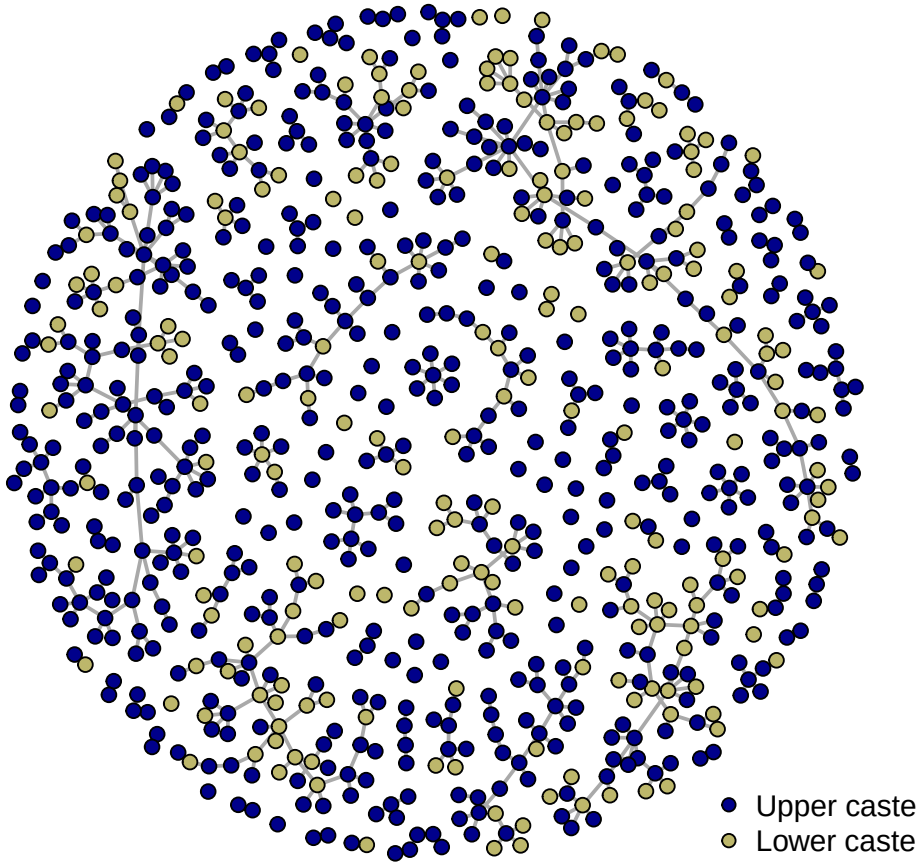
Figure B.5: Descriptive comparison of worker caste composition and industry premia



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.

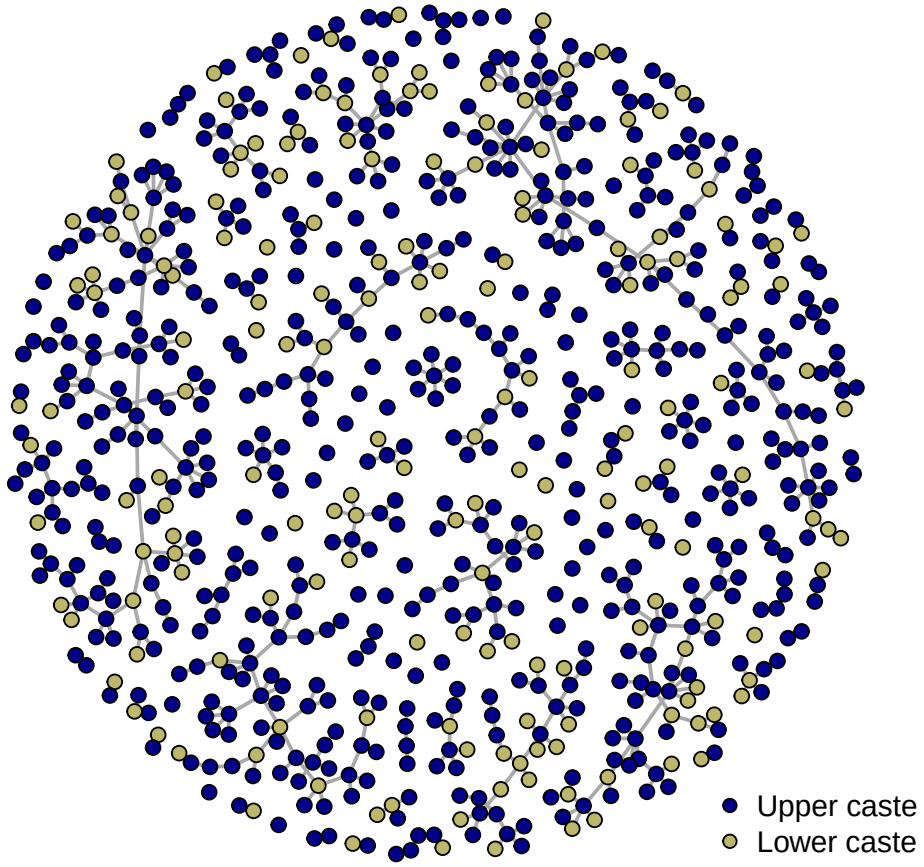
B.2 Baseline firm survey

Figure B.6: Incumbent worker network at baseline



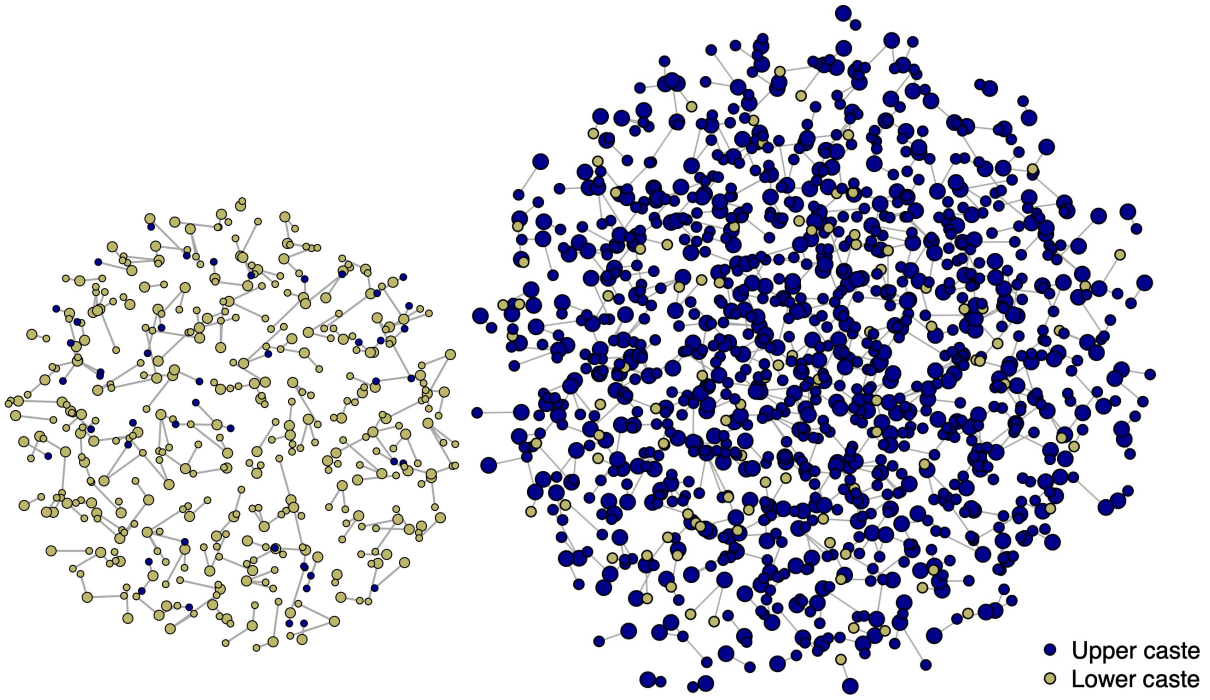
This figure represents a network graph by caste for incumbent workers at the firm at the time of the baseline, with links between referrers and referrals.

Figure B.7: Incumbent worker network at baseline (if random)



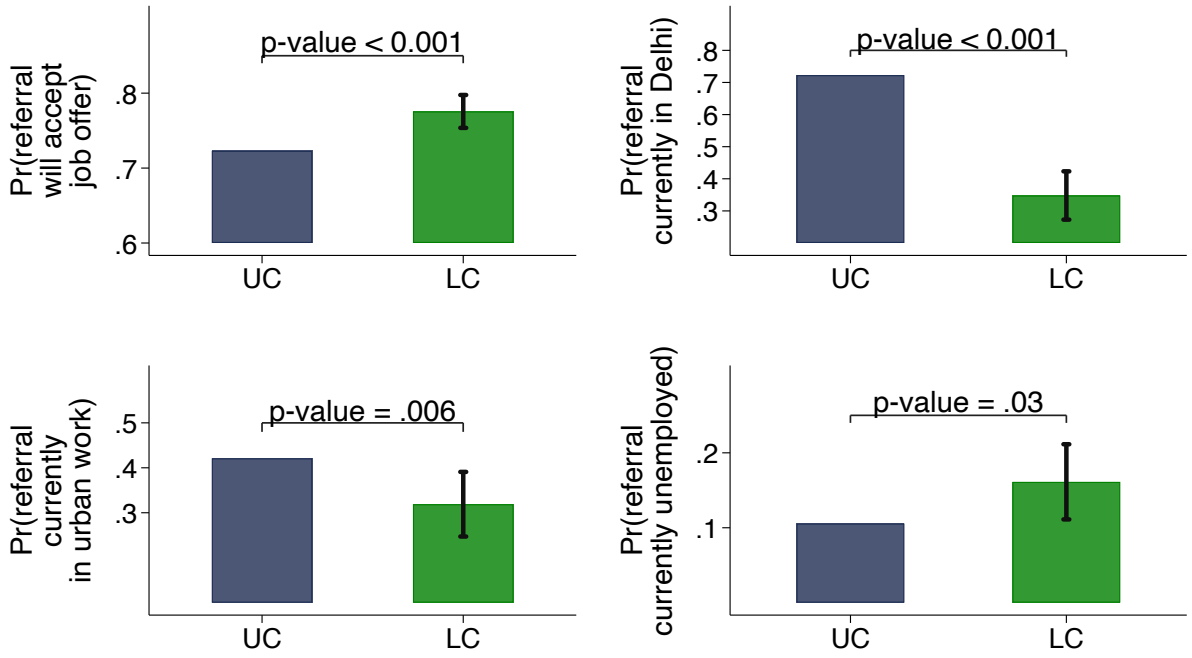
This figure represents a benchmark network graph relative to the previous slide, representing what links would look like if referrals were made proportionately by caste based on their share in the workforce.

Figure B.8: Network of referral elicitation



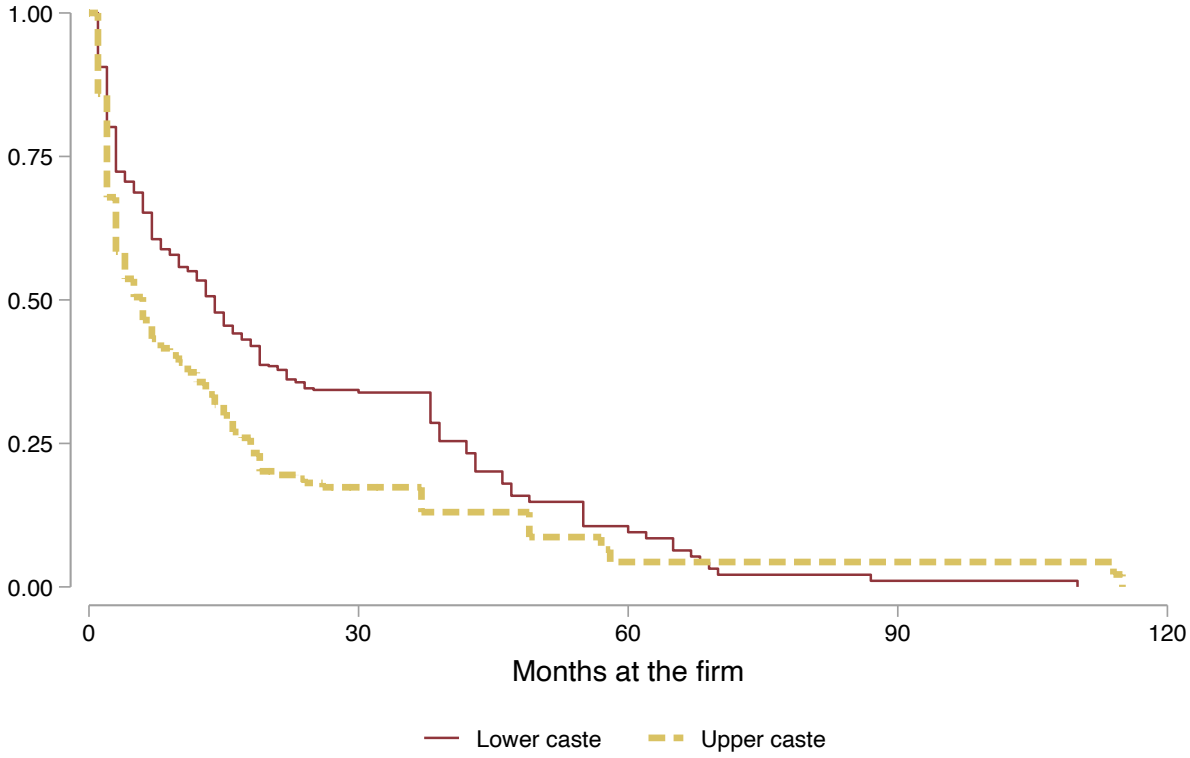
This figure represents a network graph focusing on the links between incumbent workers and their referral candidates, elicited at baseline. The graphs are represented separately for UC and LC incumbent workers, with each incumbent worker listing an average of 2.2 referral candidates.

Figure B.9: Descriptive comparison of referral candidates (firm experiment)



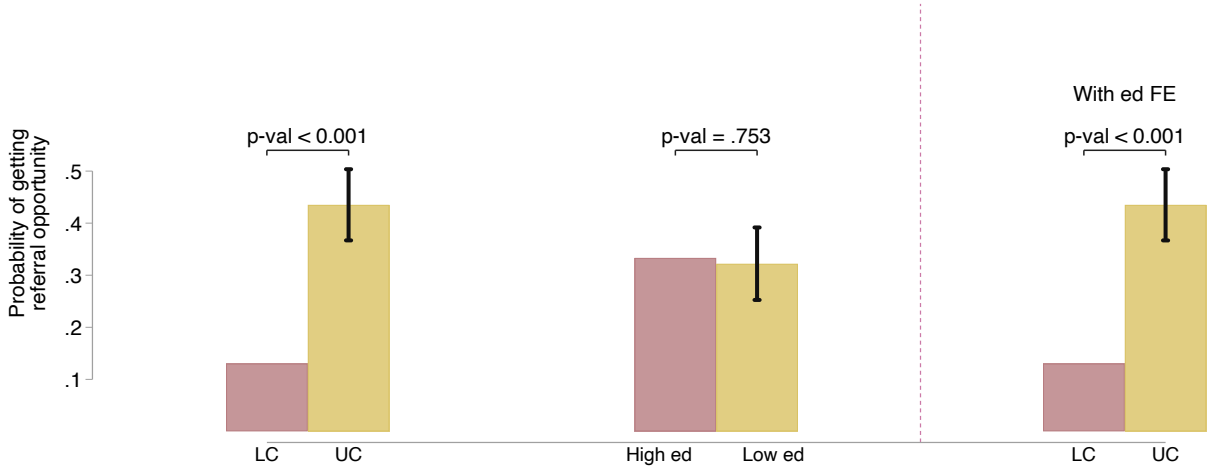
In this figure, I report simple descriptive comparisons of the referral candidates of lower caste and upper caste incumbent workers at the firm, as reported by the referrers (i.e. the incumbent workers) themselves during the baseline survey. The first graph reports the difference in lower and upper caste 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.

Figure B.10: 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.

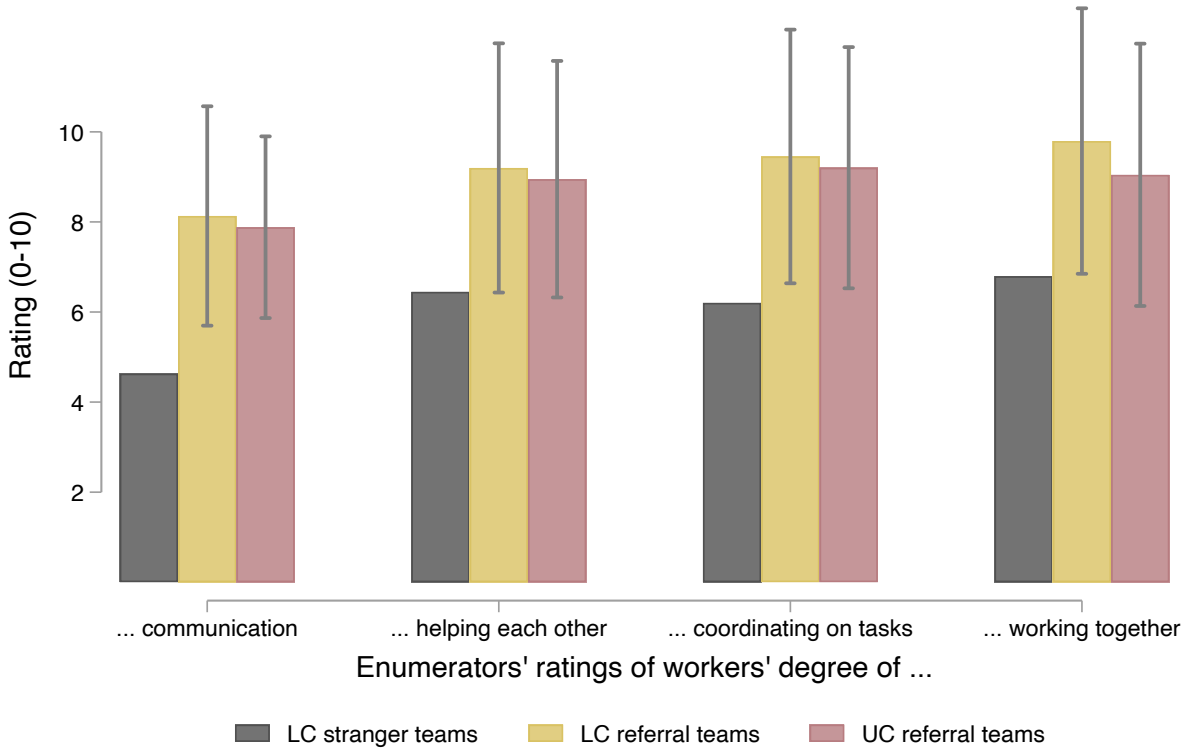
Figure B.11: Targeting of referral opportunities at baseline



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.

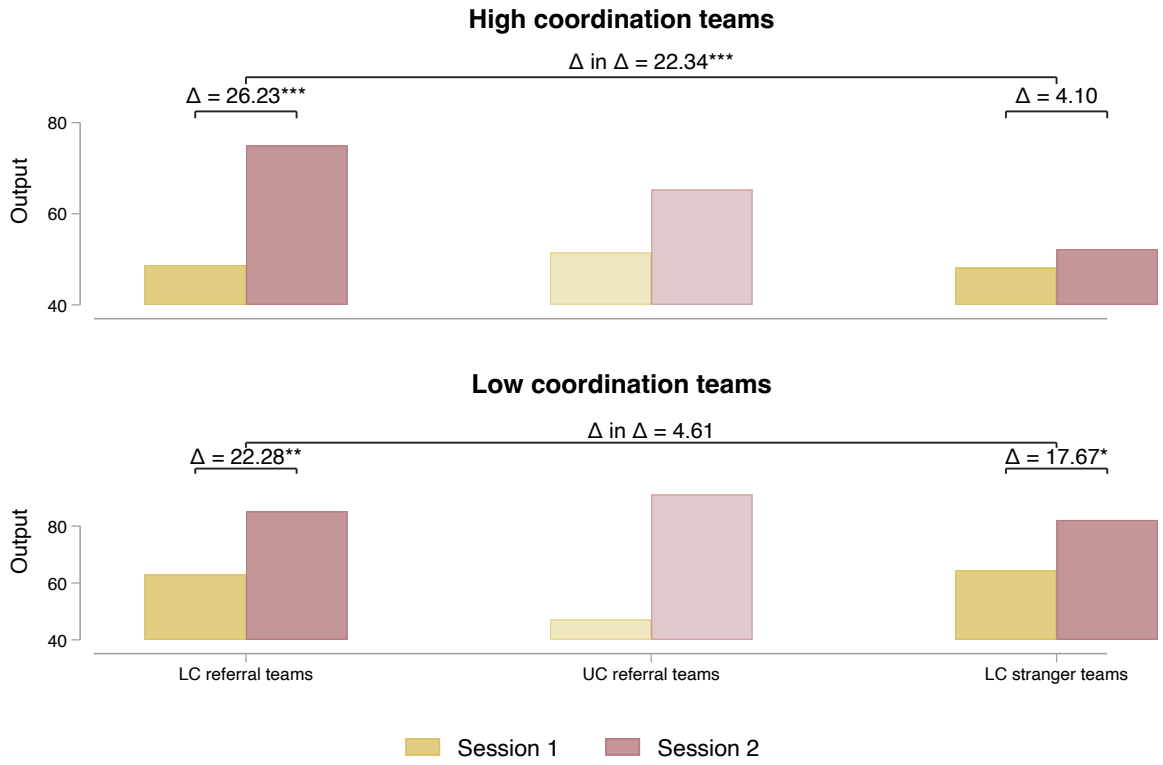
B.3 Lab-in-field survey

Figure B.12: 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.

Figure B.13: 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 B.1: 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 strangers 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 B.3: 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 strangers 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. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table B.2: 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 strangers 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 B.4: 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 strangers 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. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table B.5: 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 strangers 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. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table B.6: 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 strangers 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. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table B.7: 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 strangers 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. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

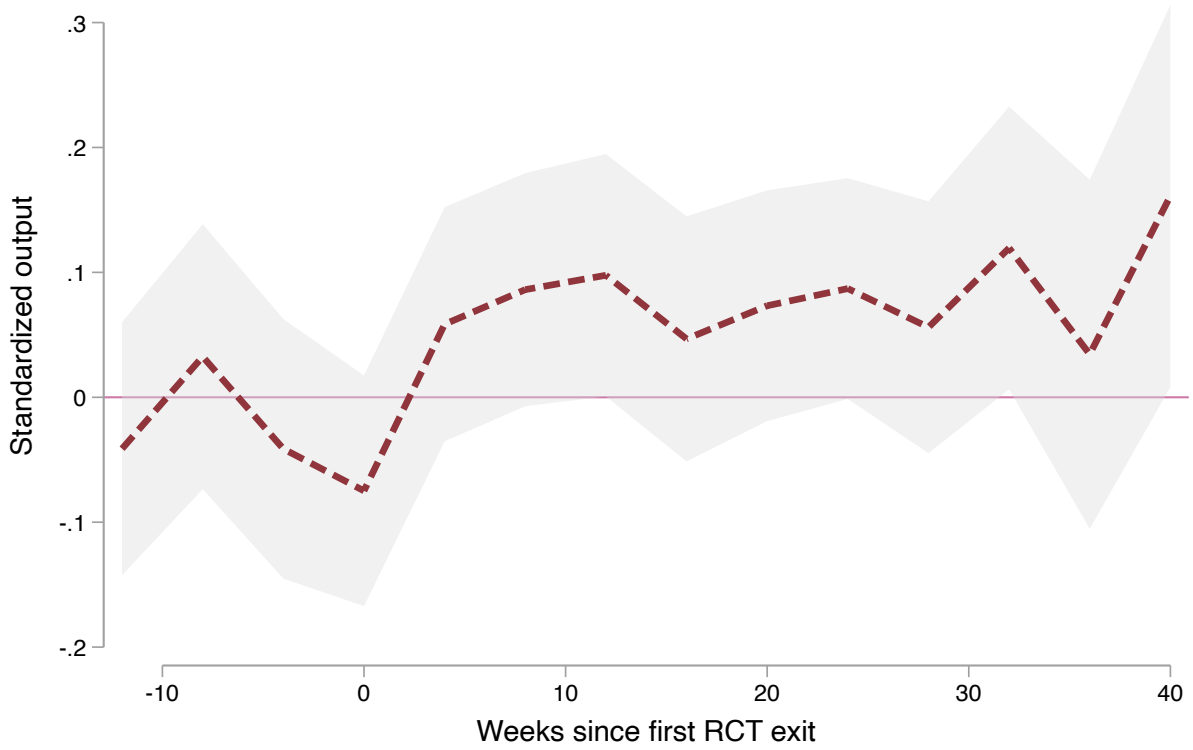
Table B.8: Lab-in-field experiment: Reason for workshop

	Workshop conducted for research		
	(1) All	(2) High coord	(3) Low coord
LC referral teams	-0.045	0.093	-0.277**
	(0.066)	(0.066)	(0.126)
UC referral teams	0.119	-0.047	0.413***
	(0.110)	(0.109)	(0.120)
LC strangers mean	0.344	0.267	0.471
R^2	0.03	0.07	0.28
Observations	470	294	176

This table reports responses to a question where workers are asked what they believe the purpose of the lab-in-field workshop was. Standard errors are clustered at the team level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

B.4 Treatment effects on firm administrative data

Figure B.14: 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 B.9: 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. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table B.10: Heterogeneous treatment effects on productivity

	Output					
Treatment	0.079	0.062	0.129*	0.088**	0.144**	0.146**
	(0.058)	(0.056)	(0.068)	(0.040)	(0.059)	(0.064)
High baseline churn	0.064					
	(0.059)					
Treatment × High baseline churn	0.010					
	(0.081)					
High baseline LC share		-0.227				
		(0.323)				
Treatment × High baseline LC share		0.046				
		(0.079)				
High baseline cohesion			0.062			
			(0.051)			
Treatment × High baseline cohesion			-0.083			
			(0.091)			
High baseline conflict risk				-0.068		
				(0.084)		
Treatment × High baseline conflict risk				-0.050		
				(0.104)		
High baseline size					0.057	
					(0.051)	
Treatment × High baseline size					-0.130	
					(0.086)	
High baseline cohesion cost beliefs						-0.025
						(0.058)
Treatment × High baseline cohesion cost beliefs						-0.104
						(0.088)
Control mean	0.003	0.003	0.003	0.003	0.003	-0.049
R^2	0.75	0.75	0.75	0.75	0.75	0.78
Observations	5715	5715	5715	5715	5715	4590

This table reports heterogeneous treatment effects for standardized productivity, with the basic specification wherein each variable that I study heterogeneity on defined as a binary for above or below median for each of the continuous variables. 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 B.11: Heterogeneous treatment effects on churn

	Churn					
Treatment	-0.938***	-0.950***	-0.974***	-1.186***	-1.087***	-1.222***
	(0.166)	(0.159)	(0.175)	(0.127)	(0.161)	(0.195)
High baseline churn	0.267					
	(0.207)					
Treatment × High baseline churn	-0.425					
	(0.263)					
High baseline LC share		1.394				
		(1.102)				
Treatment × High baseline LC share		-0.408				
		(0.250)				
High baseline cohesion			0.092			
			(0.212)			
Treatment × High baseline cohesion			-0.378			
			(0.274)			
High baseline conflict risk				-0.855**		
				(0.363)		
Treatment × High baseline conflict risk				0.918		
				(0.741)		
High baseline size					0.347	
					(0.220)	
Treatment × High baseline size					-0.170	
					(0.258)	
High baseline cohesion cost beliefs						-0.001
						(0.243)
Treatment × High baseline cohesion cost beliefs						0.065
						(0.292)
Control mean	2.785	2.785	2.785	2.785	2.785	2.798
R^2	0.16	0.16	0.16	0.16	0.16	0.16
Observations	889	889	889	889	889	714

This table reports heterogeneous treatment effects for the sum of exits and entries in each team per month, with the basic specification wherein each variable that I study heterogeneity on defined as a binary for above or below median for each of the continuous variables. All standard errors are clustered at the team level, and the regressions inlcude fixed effects at the randomization strata and month level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

B.5 Treatment effects on endline survey at firm

Table B.12: 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

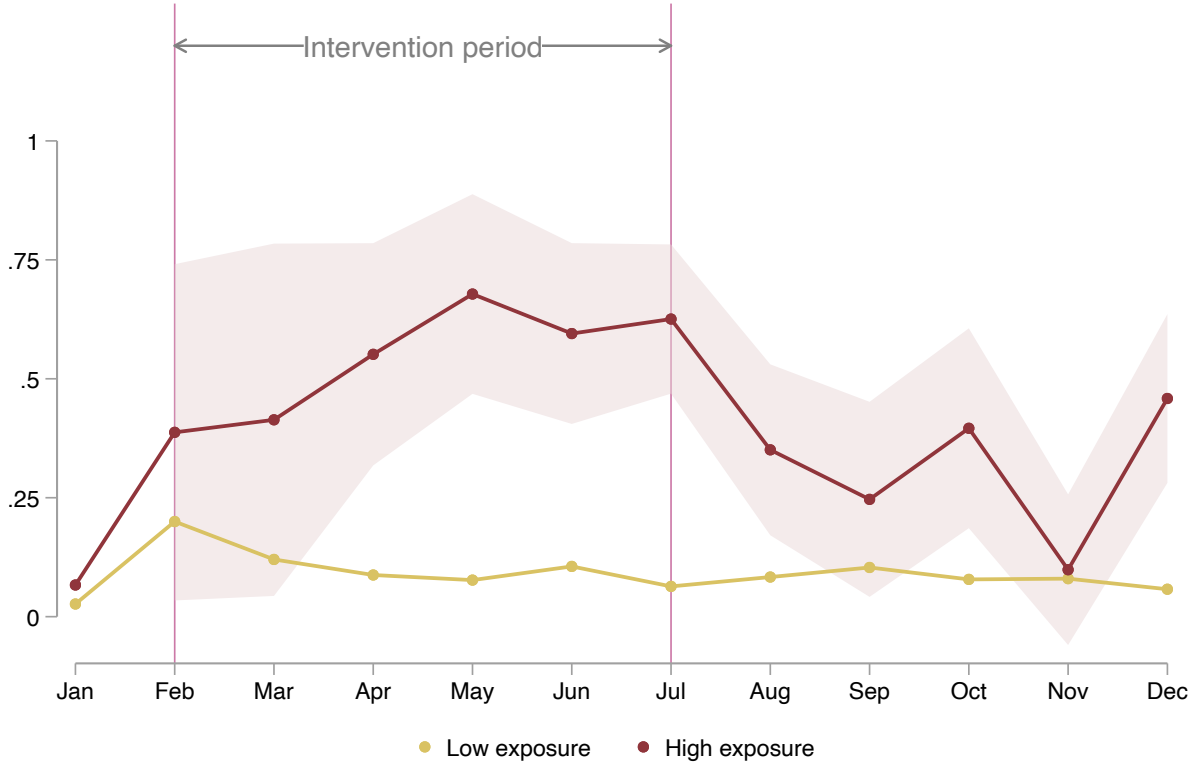
This table reports heterogeneous treatment effects on self reported cohesion, with respect to initial team diversity and the caste group of the respondents. All standard errors are clustered at the team level, and the regressions inlcude fixed effects at the randomization strata level. $*$ $p < 0.1$,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

Table B.13: 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. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Figure B.15: Referral allocation in control teams by LC exposure



This figure plots the evolution of referral allocations within the control group of teams, with each line plot depicting the subset of control teams with randomly induced above or below median exposure to lower caste referral candidates.

Table B.14: 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. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

B.6 Treatment effects on referral candidate follow-ups

Table B.19: Referral candidates' earnings

	Pr(earns anything)		Ln(1 + earnings)	
	(1) Upper caste	(2) Lower caste	(3) Upper caste	(4) Lower caste
Treatment	-0.010	0.070*	-0.037	0.288
	(0.027)	(0.041)	(0.254)	(0.374)
UC vs. LC p-value		0.111		0.475
Control mean	0.919	0.870	8.447	8.021
R^2	0.00	0.01	0.00	0.00
Observations	418	200	418	200

This table reports treatment effects on the earnings and unemployment spells, i.e. connections of the incumbent workers, as reported in a phone-based follow-up survey. 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 B.15: 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 B.20: Referral candidates' migration activity

	Strongly willing to migrate		Max search distance (km)	
	(1) Upper caste	(2) Lower caste	(3) Upper caste	(4) Lower caste
Treatment	0.005	-0.000	22.656	-21.590
	(0.044)	(0.056)	(136.315)	(64.378)
UC vs. LC p-value		0.947		0.768
Control mean	0.732	0.820	791.124	731.140
R^2	0.00	-0.00	0.00	0.00
Observations	418	200	418	200

This table reports treatment effects on migration willingness and job search activity, i.e. connections of the incumbent workers, as reported in a phone-based follow-up survey. 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 B.16: 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 B.17: 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 B.18: Referral candidates' job arrival rate

	# job offers in last 6 months		Pr(any job offer from a factory)	
	(1) Upper caste	(2) Lower caste	(3) Upper caste	(4) Lower caste
Treatment	-0.174	0.418**	-0.048	0.110
	(0.157)	(0.202)	(0.051)	(0.073)
UC vs. LC p-value		0.019		0.078
Control mean	1.349	0.520	0.502	0.420
Specification	Poisson	Poisson	OLS	OLS
R^2			0.00	0.01
Observations	418	200	418	200

This table reports treatment effects on the labor market activity among referral candidates, i.e. connections of the incumbent workers, as reported in a phone-based follow-up survey. All standard errors are clustered at the team level, and the regressions inlcude fixed effects at the randomization strata level. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

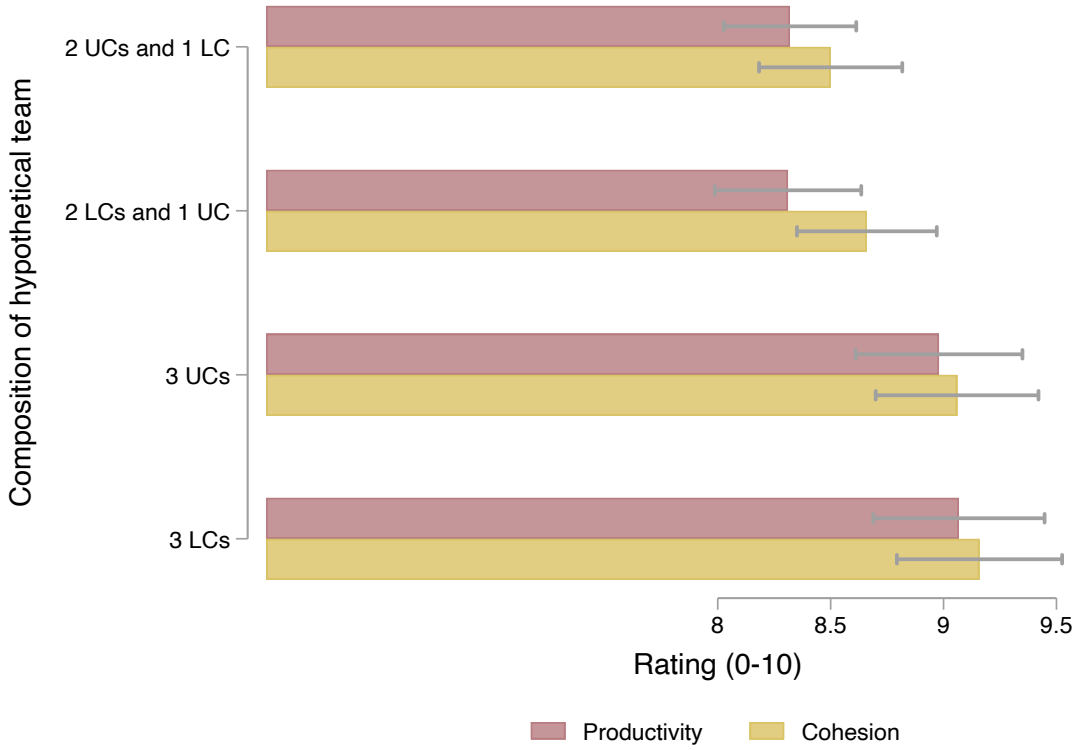
Table B.21: Vignettes survey: changes in cohesion

		Change in expected cohesion after recruitment						Obs
		Mean	Pairwise test [row - column]					
Recruitment type ↓→	Caste label ↓→		UC stranger No	UC referral No	LC stranger No	LC referral No	LC stranger Yes	
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
UC stranger	No	0.76	-					93
UC referral	No	0.93	0.16**	-				107
			(0.06)					
LC stranger	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 stranger	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. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

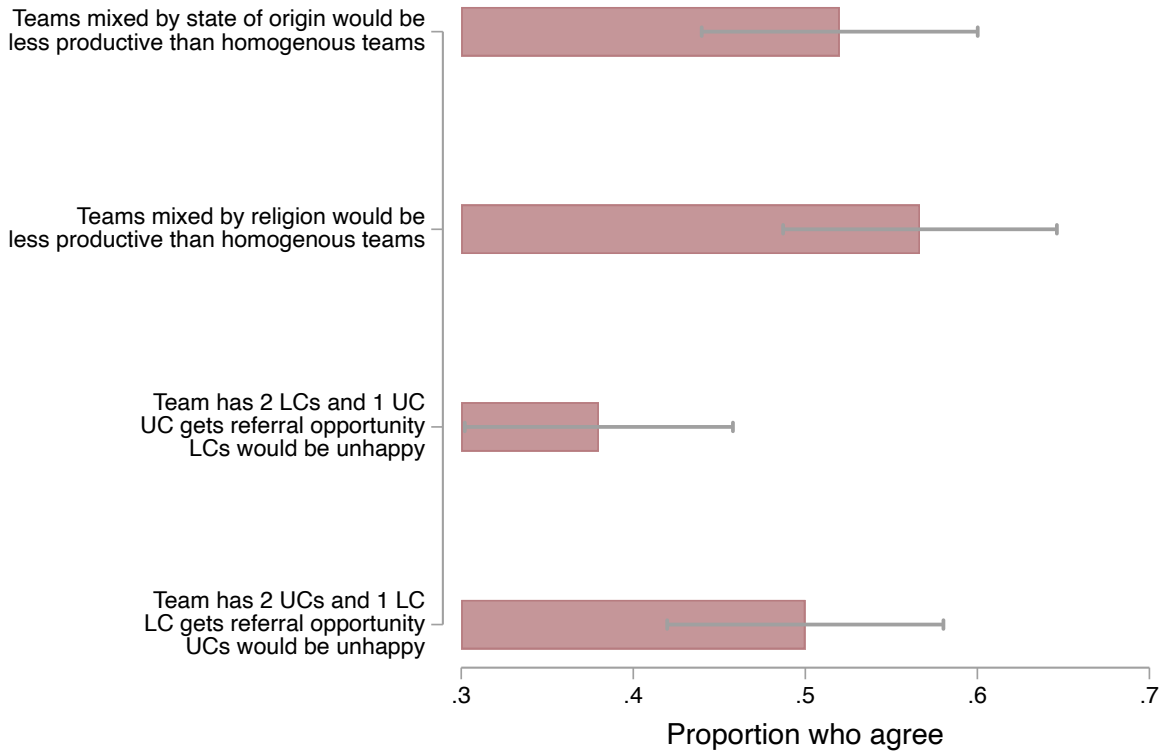
B.7 Vignettes survey

Figure B.16: Vignettes survey: expectations of productivity and cohesion vs. team composition



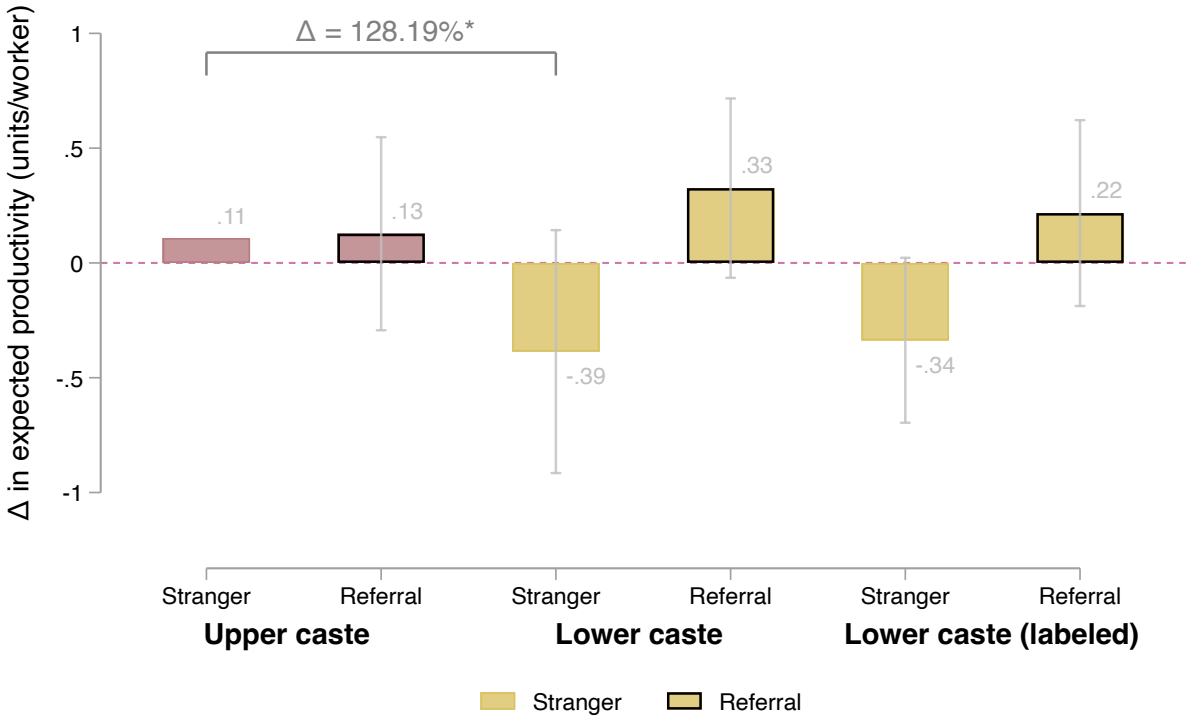
This figure summarizes responses from the vignettes survey on the expected productivity of hypothetical teams based on their composition.

Figure B.17: Vignettes survey: policy perceptions



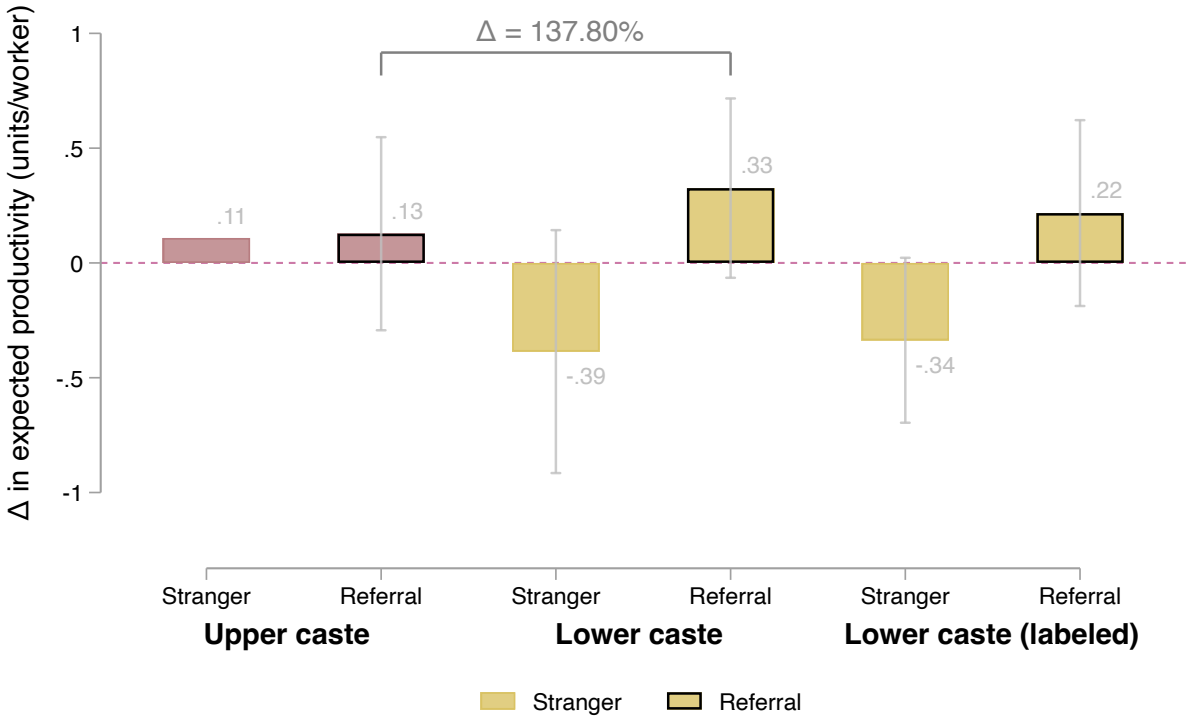
This figure summarizes responses from the vignettes survey on agreement rates to a range of statements concerning team diversity, cohesion and productivity.

Figure B.18: Vignettes survey



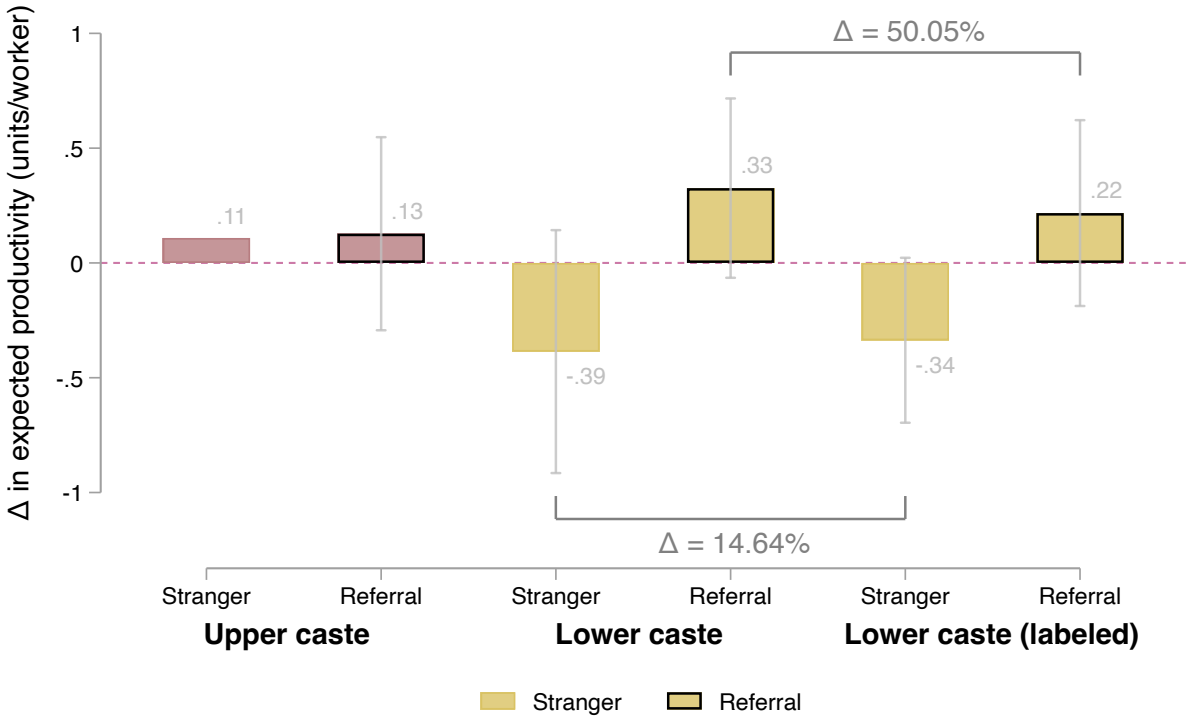
This figure plots treatment effects from the vignettes survey, focusing on pairwise comparisons. The point estimates correspond to the figures in Table 8.

Figure B.19: Vignettes survey



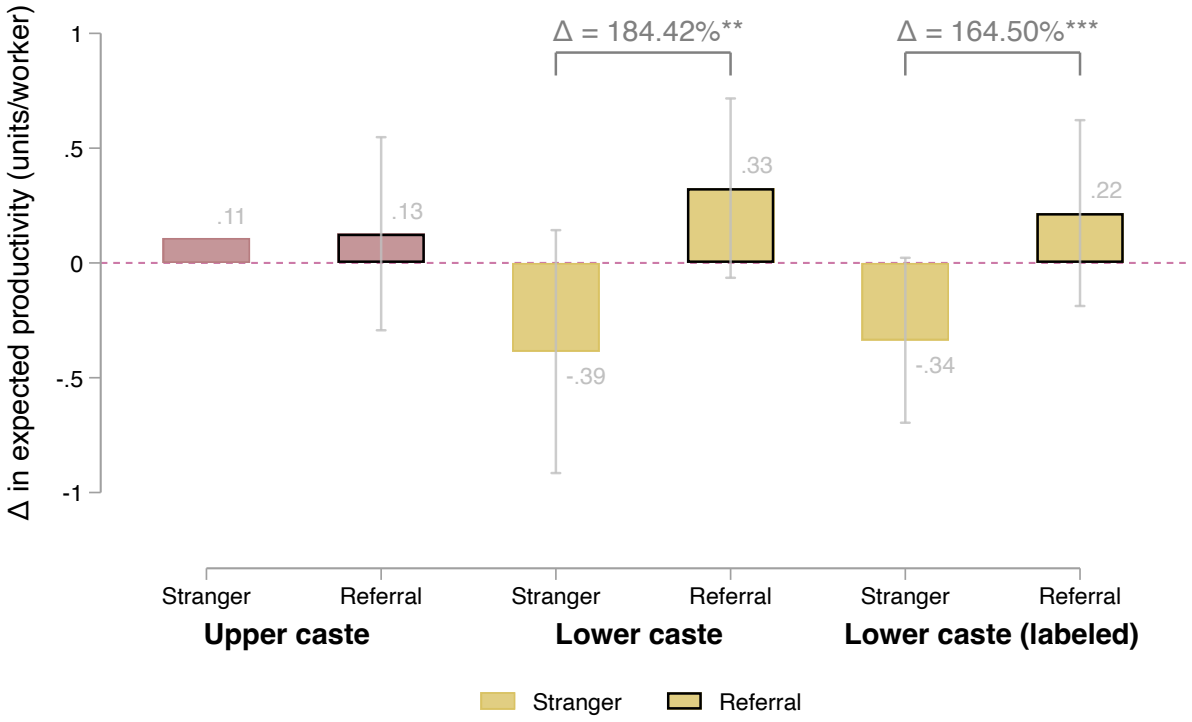
This figure plots treatment effects from the vignettes survey, focusing on pairwise comparisons. The point estimates correspond to the figures in Table 8.

Figure B.20: Vignettes survey



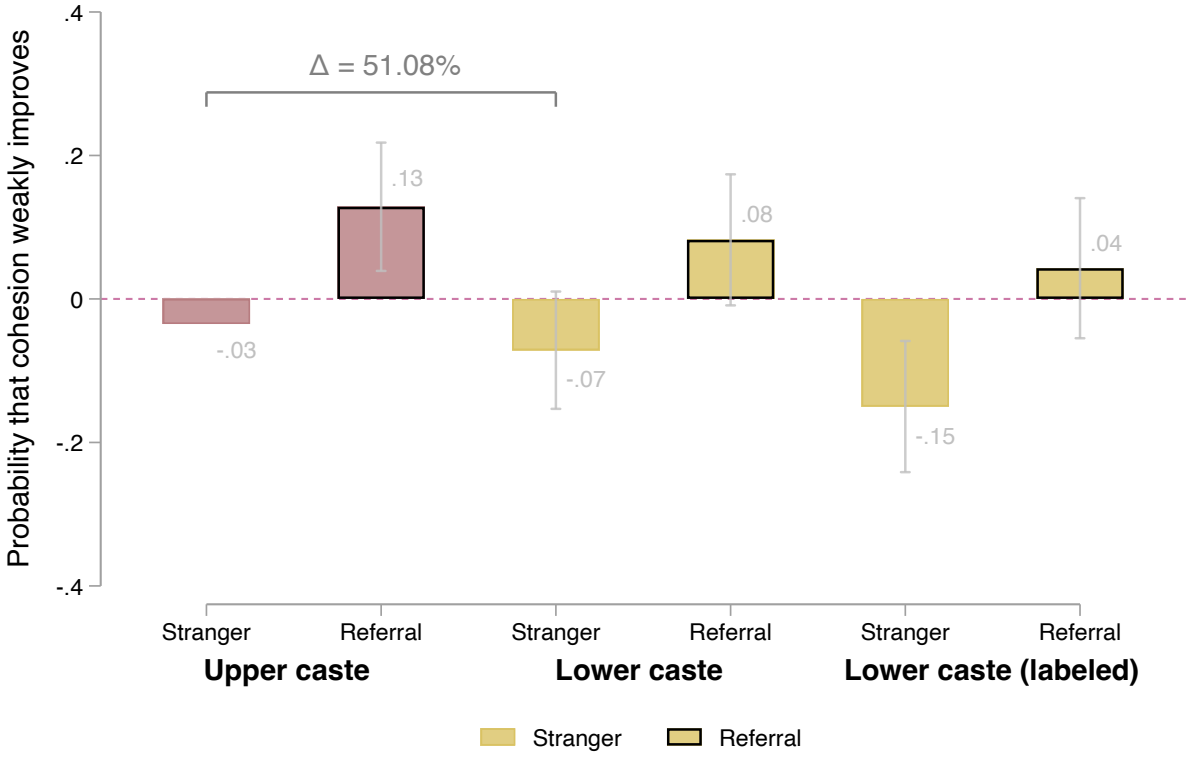
This figure plots treatment effects from the vignettes survey, focusing on pairwise comparisons. The point estimates correspond to the figures in Table 8.

Figure B.21: Vignettes survey



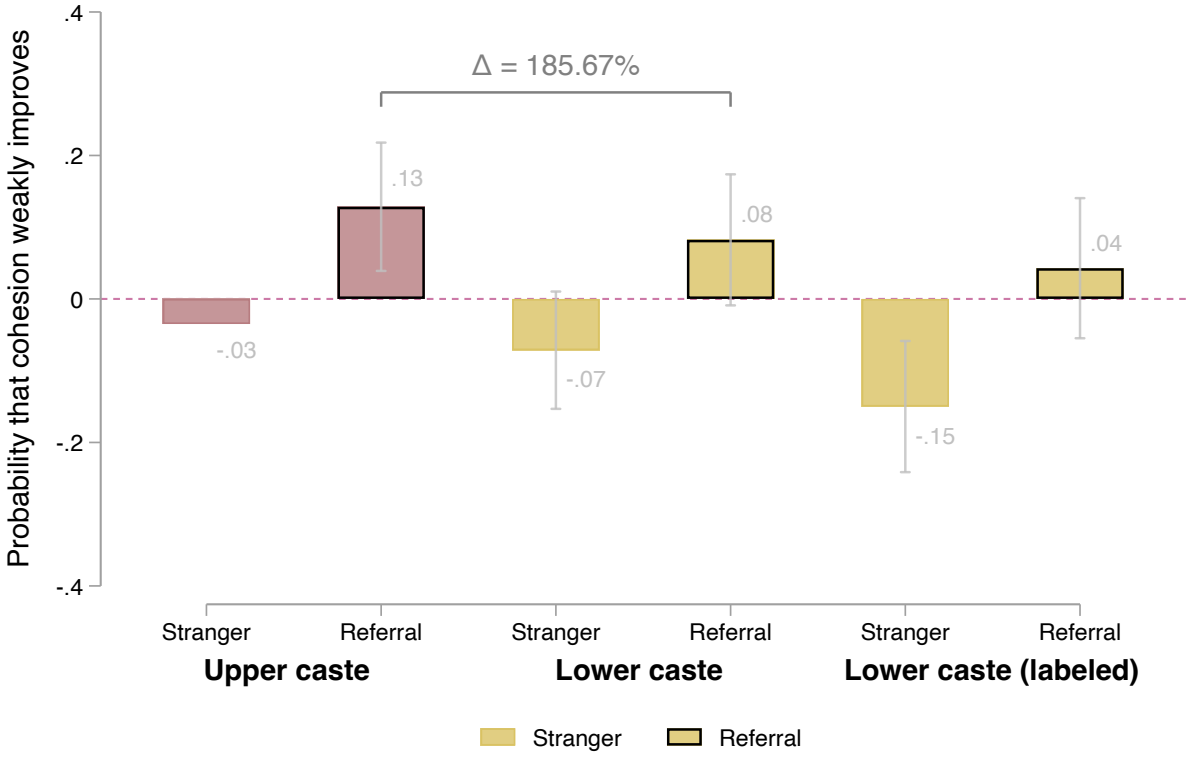
This figure plots treatment effects from the vignettes survey, focusing on pairwise comparisons. The point estimates correspond to the figures in Table 8.

Figure B.22: Vignettes survey



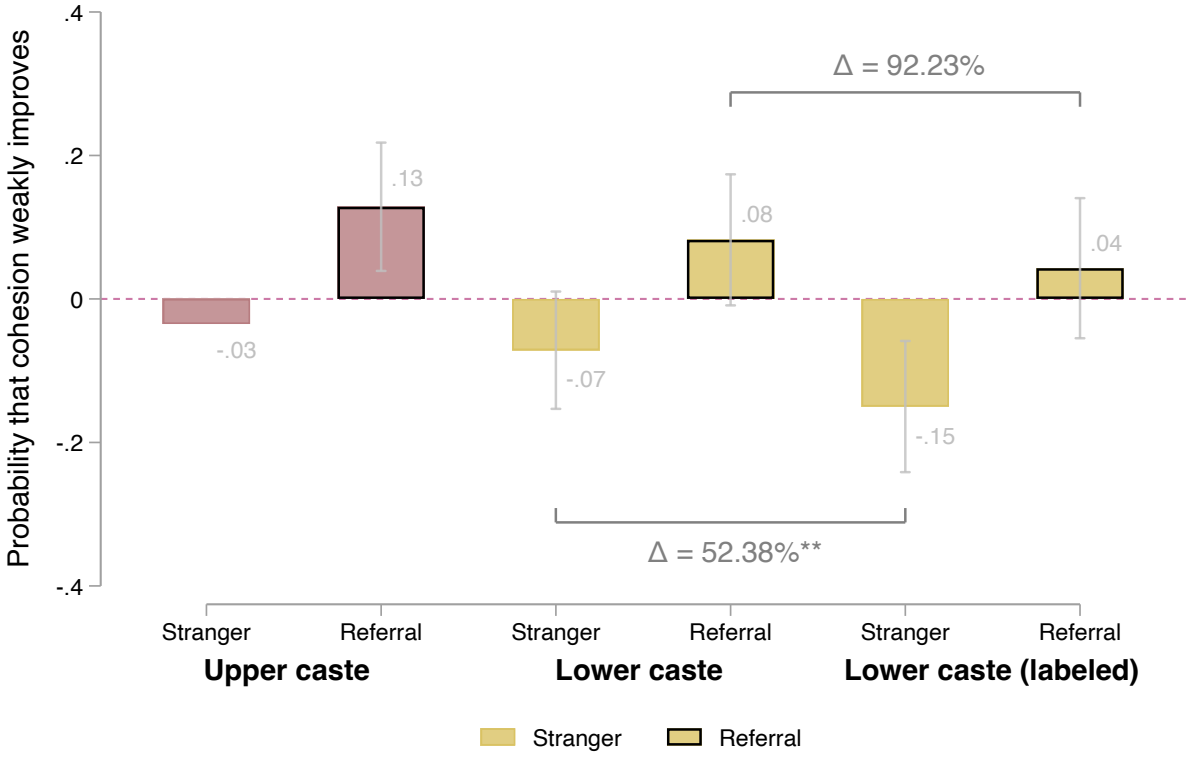
This figure plots treatment effects from the vignettes survey, focusing on pairwise comparisons. The point estimates correspond to the figures in Table B.21.

Figure B.23: Vignettes survey



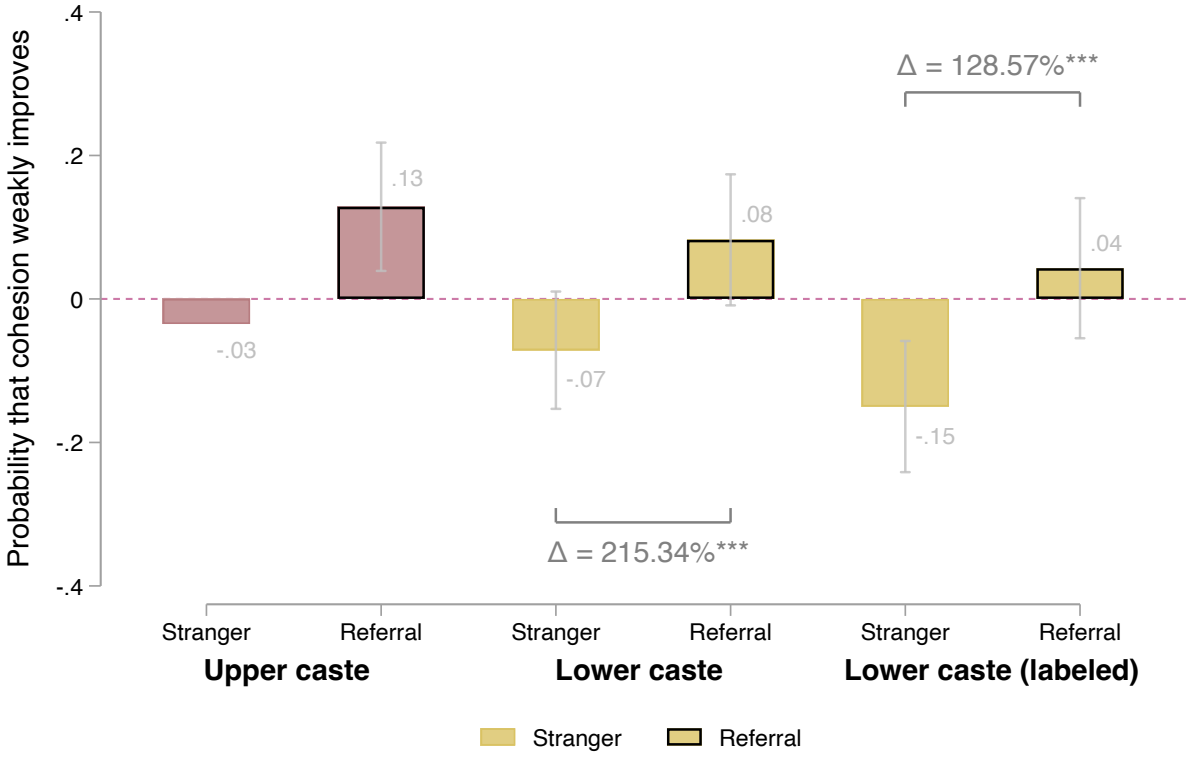
This figure plots treatment effects from the vignettes survey, focusing on pairwise comparisons. The point estimates correspond to the figures in Table B.21.

Figure B.24: Vignettes survey



This figure plots treatment effects from the vignettes survey, focusing on pairwise comparisons. The point estimates correspond to the figures in Table B.21.

Figure B.25: Vignettes survey



This figure plots treatment effects from the vignettes survey, focusing on pairwise comparisons. The point estimates correspond to the figures in Table B.21.

Table B.22: Effect on expected output (robustness)

Recruitment type ↓→	Caste label ↓→	Change in expected output per worker after recruitment						Obs
		Mean	Pairwise test [row - column]					
			UC stranger No	UC referral No	LC stranger No	LC referral No	LC stranger Yes	
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
UC stranger	No	-0.32	-					38
UC referral	No	0.44	0.76*	-				62
			(0.45)					
LC stranger	No	-0.31	0.00	-0.75	-			55
			(0.58)	(0.51)				
LC referral	No	0.58	0.89*	0.14	0.89	-		45
			(0.51)	(0.43)	(0.57)			
LC stranger	Yes	0.17	0.48	-0.27	0.48	-0.41	-	57
			(0.48)	(0.39)	(0.54)	(0.46)		
LC referral	Yes	0.15	0.46	-0.29	0.46	-0.43	-0.02	43
			(0.46)	(0.37)	(0.52)	(0.44)	(0.40)	

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. This table only keeps the first policy that is offered to the respondent, as a robustness check to the main specification reported in Table 8. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table B.23: Effect on expected cohesion (robustness)

Recruitment type ↓→	Caste label ↓→	Change in expected cohesion after recruitment						Obs
		Mean	Pairwise test [row - column]					
			UC stranger No	UC referral No	LC stranger No	LC referral No	LC stranger Yes	
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
UC stranger	No	0.74	-					38
UC referral	No	0.92	0.18**	-				62
			(0.08)					
LC stranger	No	0.76	0.03	-0.16**	-			55
			(0.09)	(0.07)				
LC referral	No	0.91	0.17**	-0.01	0.15**	-		45
			(0.08)	(0.06)	(0.07)			
LC stranger	Yes	0.72	-0.02	-0.20***	-0.04	-0.19**	-	57
			(0.09)	(0.07)	(0.08)	(0.07)		
LC referral	Yes	0.86	0.12	-0.06	0.10	-0.05	0.14*	43
			(0.09)	(0.06)	(0.08)	(0.07)	(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. This table only keeps the first policy that is offered to the respondent, as a robustness check to the main specification reported in Table B.21. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

C Robustness tests

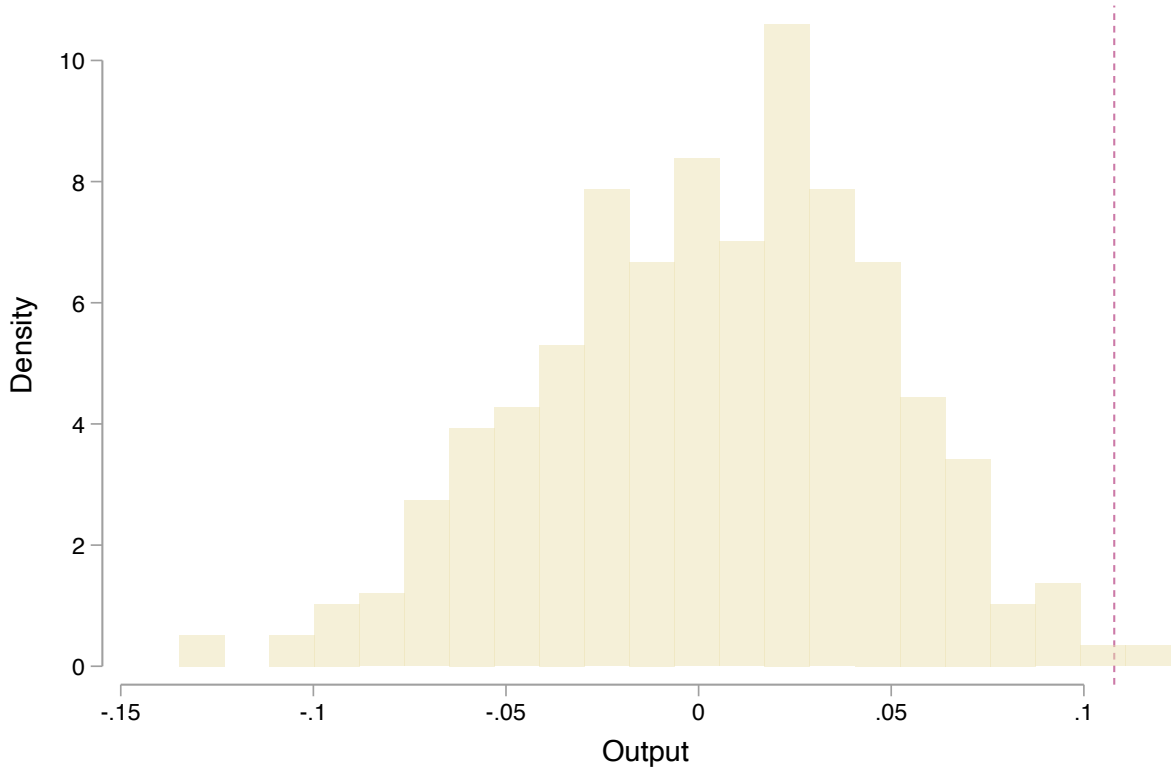
C.1 Randomization inference

Table C.1: Summary of randomization inference p-values

Variable	Standard p-value (1)	Randomization inference p-value (2)	Table in main text (3)
Output (pooled)	0.011	0.060	Table 3, Column 1
Output (high coordination)	0.004	0.076	Table 3, Column 3
Output (low coordination)	0.239	0.074	Table 3, Column 5
Team size	0.136	0.082	
Incumbent share	0.156	0.072	
LC share	0.000	0.066	
Total exits	0.000	0.054	Table 4, Column 1
Incumbent exits	0.198	0.060	Table 4, Column 2
New hire exits	0.000	0.076	Table 4, Column 3
LC exits	0.060	0.060	Table 4, Column 4
Total entries	0.000	0.074	Table 4, Column 5
LC entries	0.000	0.056	Table 4, Column 6
Total churn	0.000	0.066	Table 4, Column 7
Referrals to LCs	0.000	0.080	Table 2, Column 1
Referrals to LCs / Total referrals	0.000	0.072	Table 2, Column 2

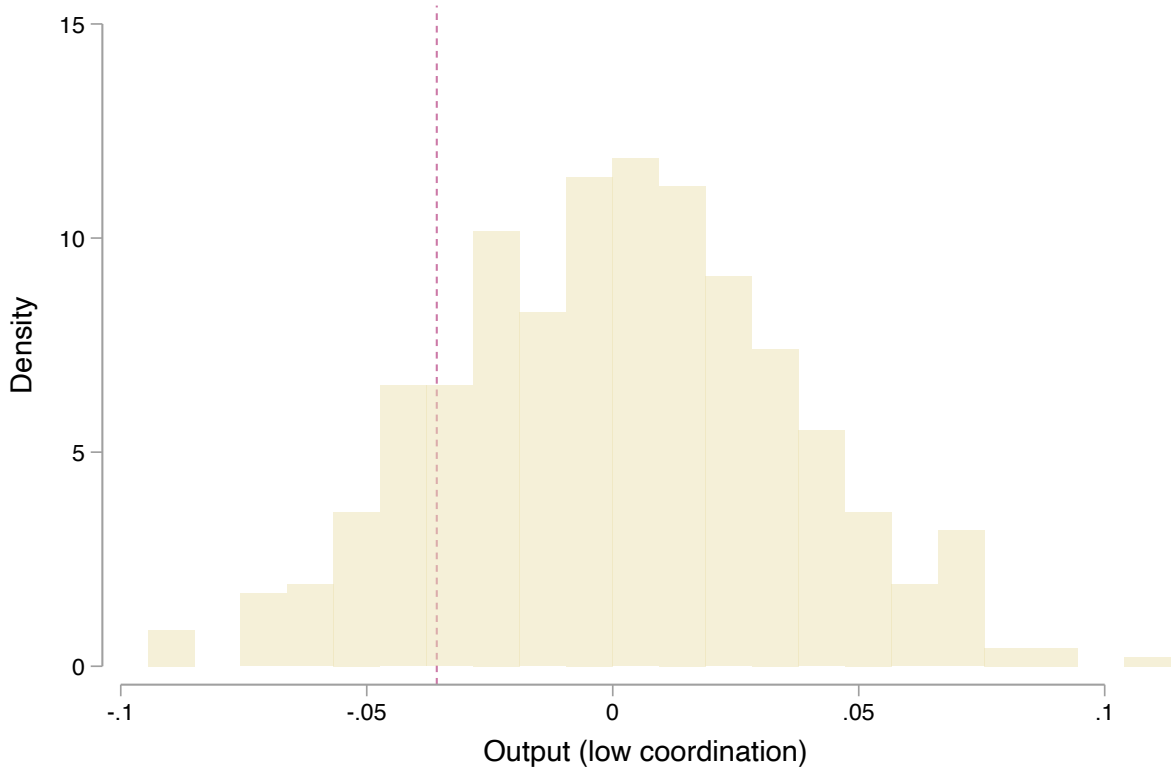
This table presents a summary of p-values from randomization inference from $N = 500$ simulations. For each given outcome variable, column 1 lists the standard p-value, column 2 lists the randomization inference p-value, with labels corresponding to tables in the main text in column 3.

Figure C.1: RI estimates of treatment effects on output



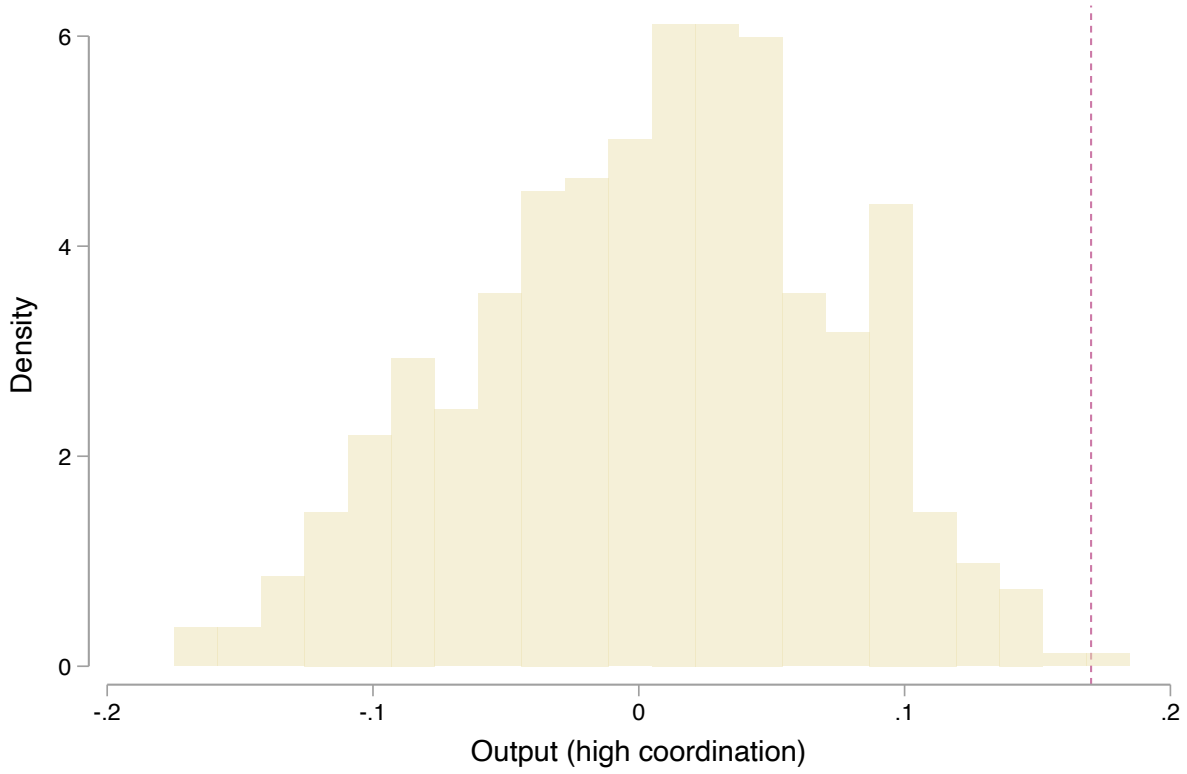
This figure plots the distribution of p-values from randomization inference simulations for the output listed in the figure caption, with the vertical dotted line denoting the standard p-value reported in the main text.

Figure C.2: RI estimates of treatment effects on output (low coordination teams)



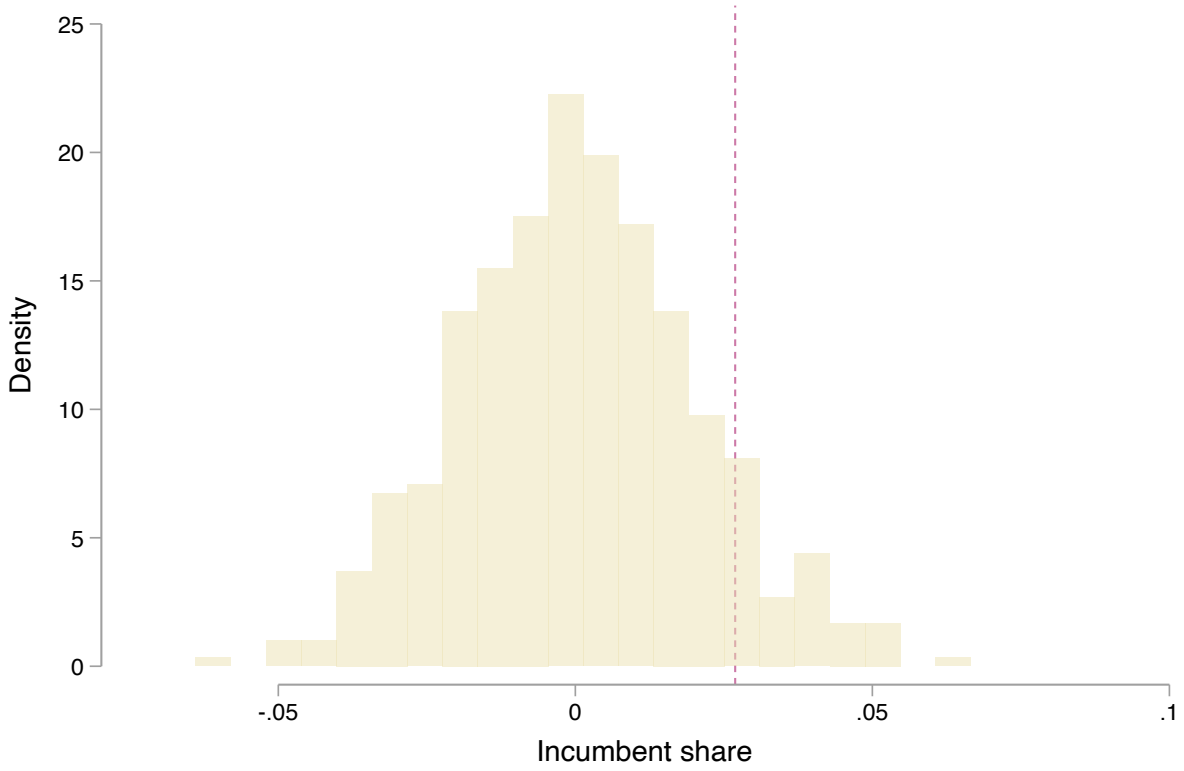
This figure plots the distribution of p-values from randomization inference simulations for the output listed in the figure caption, with the vertical dotted line denoting the standard p-value reported in the main text.

Figure C.3: RI estimates of treatment effects on output (high coordination teams)



This figure plots the distribution of p-values from randomization inference simulations for the output listed in the figure caption, with the vertical dotted line denoting the standard p-value reported in the main text.

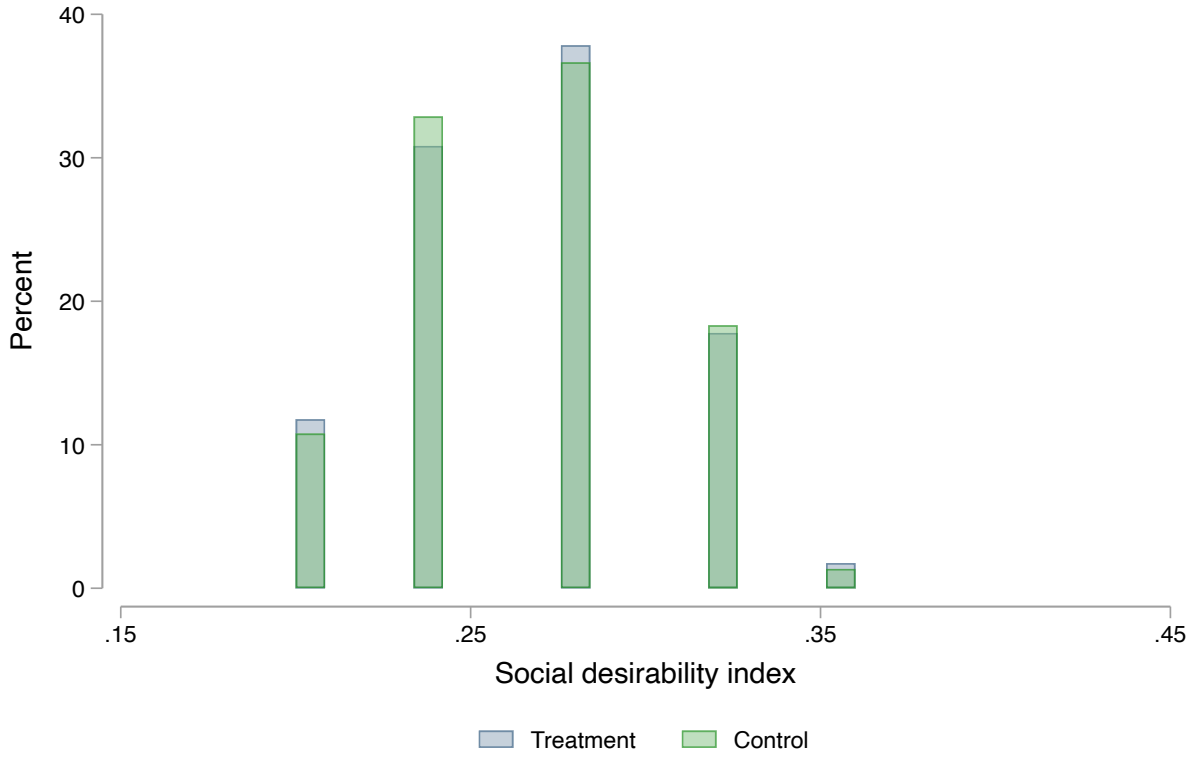
Figure C.4: RI estimates of treatment effects on team LC shares



This figure plots the distribution of p-values from randomization inference simulations for the output listed in the figure caption, with the vertical dotted line denoting the standard p-value reported in the main text.

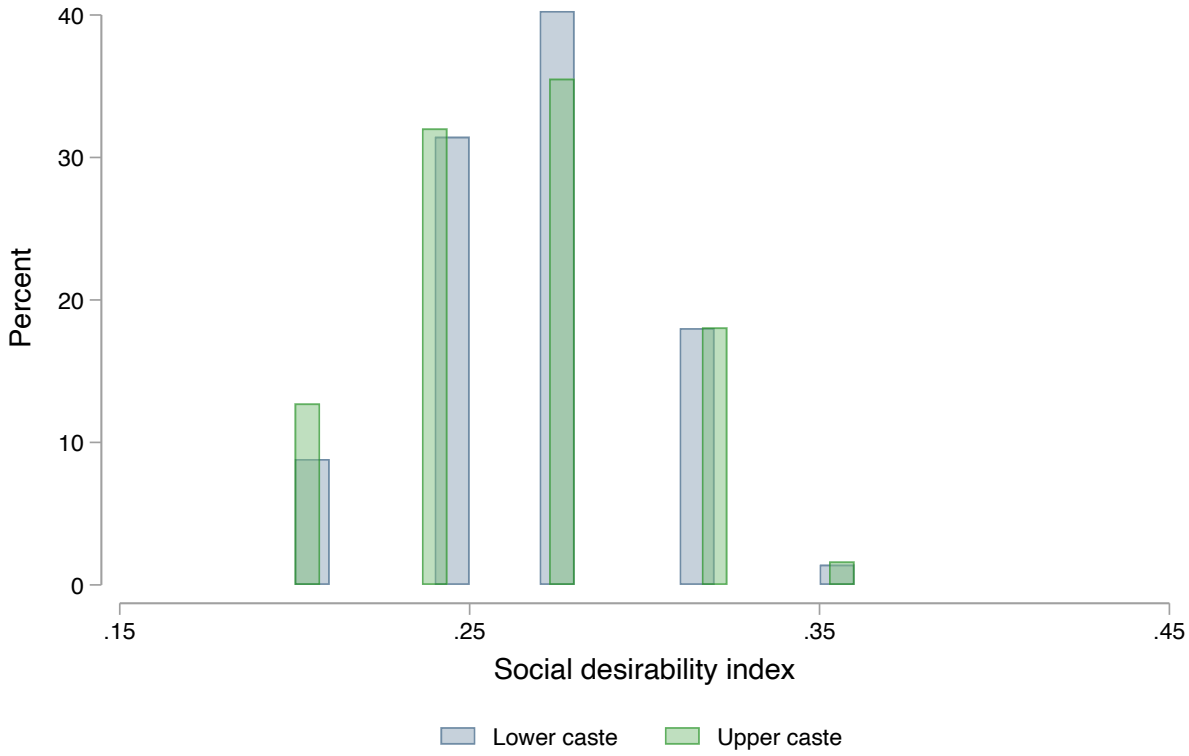
C.2 Social desirability bias

Figure C.5: Distribution of social desirability index by treatment



This figure plots the distribution of a social desirability index created from responses to questions from an adapted version of a Marlowe-Crowne Social Desirability Scale, reported separately for workers at the firm corresponding to treatment and control teams.

Figure C.6: Distribution of social desirability index by caste status



This figure plots the distribution of a social desirability index created from responses to questions from an adapted version of a Marlowe-Crowne Social Desirability Scale, reported separately for workers at the firm corresponding to lower caste and upper caste workers.

Table C.2: Heterogeneity by social desirability

	Team cohesion (0-10)		Friends outside		Bargaining power		Value of referral		Job satisfaction	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment	0.098	-0.246*	0.209	0.031	-0.026	-0.028	0.005	-0.035	0.161	0.189
	(0.092)	(0.147)	(0.403)	(0.552)	(0.039)	(0.048)	(0.039)	(0.048)	(0.116)	(0.169)
Above median social desirability (SD)	-0.208	-0.406*	-1.354***	-1.512***	-0.116*	-0.112	-0.061	-0.035	0.360**	0.315*
	(0.202)	(0.218)	(0.370)	(0.434)	(0.065)	(0.072)	(0.062)	(0.075)	(0.165)	(0.183)
Treatment × Above median SD	0.291	0.549*	1.166	-0.051	0.011	-0.046	-0.023	-0.010	-0.294	-0.499
	(0.259)	(0.322)	(0.920)	(0.733)	(0.089)	(0.112)	(0.080)	(0.102)	(0.248)	(0.356)
Lower caste		0.581***		-0.896		-0.038		-0.096		-0.159
		(0.156)		(0.558)		(0.065)		(0.076)		(0.213)
Treatment × Lower caste		0.301		0.890		0.027		0.136		0.043
		(0.228)		(0.899)		(0.077)		(0.088)		(0.289)
Lower caste × Above median SD		0.663*		0.778		-0.010		-0.090		0.207
		(0.367)		(0.838)		(0.134)		(0.154)		(0.498)
Treatment × Lower caste × Above median SD		-0.696		2.122		0.121		0.010		0.307
		(0.476)		(1.867)		(0.179)		(0.188)		(0.631)
Control mean	8.312	8.312	3.485	3.485	0.679	0.679	0.739	0.739	8.295	8.295
R^2	0.02	0.10	0.02	0.03	0.02	0.02	0.02	0.03	0.02	0.02
Observations	767	767	767	767	770	770	770	770	769	769

This table reports the primary outcomes from the endline survey at the firm, with the treatment indicator interacted with whether the respondent was above or below the median in the social desirability index, along with an interaction by caste status. Outcomes are as defined in the main text. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

D Attrition and non-response

Table D.1: Sample comparison

Variable	Firm sample (F)		Vignettes sample (V)		LiF sample (L)		Pairwise test p-values			Observations
	Mean	SD	Mean	SD	Mean	SD	C vs. V	C vs. F	V vs. F	
Monthly income	11.978	4.803	0.000	0.000	12.150	9.002	.	0.000	.	
State of origin: Bihar	0.337	0.473	0.600	0.492	0.514	0.501	0.000	0.000	0.000	
State of origin: UP	0.539	0.499	0.300	0.460	0.311	0.464	0.000	0.000	0.000	
State of origin: Delhi	0.053	0.223	0.073	0.262	0.143	0.350	0.000	0.000	0.000	
Participated in cultivation last year	0.230	0.421	0.000	0.000	0.371	0.484	.	0.000	.	

This table presents a comparison of the sample of respondents across the three experiments: the firm workforce, the sample for the vignettes survey, and the sample for the lab in field experiment. In rows, I report means and standard deviations for indicators that are common across the three samples where available. I also present pairwise t-tests between each of the samples. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

Table D.2: Attrition in referral candidate survey

	(1) Not surveyed		(2) Surveyed		T-test Difference
Variable	N/[Clusters]	Mean/SE	N/[Clusters]	Mean/SE	(1)-(2)
Same caste as referrer	573	0.607	454	0.656	-0.049
	[414]	(0.027)	[409]	(0.024)	
Likelihood of accepting job	573	6.981	454	7.870	-0.889***
	[414]	(0.078)	[409]	(0.073)	
In Delhi as of the baseline	472	0.612	406	0.621	-0.008
	[352]	(0.025)	[363]	(0.024)	
Unemployed as of baseline	573	0.401	454	0.414	-0.013
	[414]	(0.024)	[409]	(0.024)	
Expected quality on the job	573	6.806	454	6.830	-0.024
	[414]	(0.084)	[409]	(0.079)	
Close or extended family member	573	0.815	454	0.804	0.011
	[414]	(0.020)	[409]	(0.020)	

This table presents a comparison of the sample of referral candidates that I was able to track in the endline follow-up survey and those that did not respond. I represent the sample size, means and standard errors of the entire set of variables I observe for the entire universe of referral candidates. Column 3 represents the difference and statistical significance from a pairwise t-test. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$

E Deviations from pre-analysis plan

This study was pre-registered at the AEA RCT Registry (AEARCTR-0013091) on March 2, 2024, with the final update to the pre-analysis plan recorded on August 24, 2024. The main experimental design elements, primary outcomes, and empirical specifications followed the pre-specified plan closely. Nonetheless, there were a small number of deviations and additions, documented here.

Analogs to primary outcomes from lab-in-field experiment. As planned, I measured the primary firm-level outcomes (team composition, productivity, retention, and cohesion) using administrative data and surveys. In the lab-in-field experiment, I added two outcome measures that were subsequently listed in the initial pre-analysis plan as an edited entry: (i) productivity effects, again disaggregated by high- and low-coordination tasks as in the firm experiment, and (ii) cooperation in a symbol-matching game, coded as the share of correct matches relative to the true number of matches. These were introduced to strengthen measurement of cohesion, which in the main firm experiment relied on self-reported survey data. Separately, I added two outcome measures from the vignette surveys that were not listed in the initial pre-analysis plan but are consistent with the conceptual framework: (i) respondents' reports of change in expected productivity, and (ii) respondents' reports of change in expected cohesion, before and after hiring workers through a range of different recruitment modes.

Heterogeneous treatment effects. The PAP pre-specified testing for heterogeneity by task type, team size, baseline lower caste share, churn rate, baseline cohesion, and supervisor beliefs. I primarily report heterogeneous treatment effects in the firm (and lab-in-field) experiment by coordination, which is defined at the task type. I report the other heterogeneities in appendix tables B.10 and B.11.

Specification refinements. The PAP proposed basic OLS specifications with randomization-strata fixed effects and team-level clustering. All core results follow this specification. For some lab-in-field and follow-up candidate outcomes, I also estimated difference-in-differences specifications or interacted treatment indicators with coordination intensity, in order to align more directly with the hypotheses motivating the extension. These refinements do not alter the interpretation of treatment effects. For the analysis of the referral candidate follow-up survey outcomes, I also employ a PDS lasso specification to select controls to improve precision.

Implementation notes. Minor deviations occurred in referral allocation due to (i) administrative errors in the HR department's implementation of the prescribed referral lists, (ii) occasional refusal of referral opportunities by selected incumbents, and (iii) rare supervisor overrides. These account for the small number of non-compliant referral opportunities reported in table 2. Separately, as noted in the edit to the original pre-analysis plan, the

cross-randomization intervention where recruited workers would be shuffled across teams was not implemented due to a lack of feasibility.

F Static model

Consider a static model in which a firm allocates a share $\rho \in [0, 1]$ of referral opportunities to minority incumbent employees (LCs), with the remainder $1 - \rho$ going to majority incumbents (UCs). The firm has fixed team size L, and current team diversity is given by $S \in [0, 1]$, the share of team members who are LCs.

Production. Output depends on team cohesion and follows:

$$Y = A \cdot C(S, \rho) \cdot L^{\alpha}, \quad \text{where } A > 0,\ \alpha \in (0, 1)$$

$C(.)$ denotes team cohesion, i.e. a measure of how well the team works together. Cohesion is highest when referrals align with the team's (exogenously determined) ethnic composition, and I model misalignment costs as:

$$C(S, \rho) = 1 - \kappa|\rho - S|^{\eta}, \quad \kappa > 0,\ \eta > 0$$

Training Cost. Turnover imposes training costs on teams, and I model them as a function of the exits. Let γ_l and γ_u be inverse turnover rates (i.e., a measure of retention) for LCs and UCs, respectively, and define the retention advantage:

$$\delta = \gamma_l - \gamma_u$$

For the purpose of the parameterization, I take γ_u as given and evaluate the response of the referral allocation ρ with respect to δ.

Expected turnover is a weighted average:

$$\bar{T}(\rho) = \rho \cdot \frac{1}{\gamma_l} + (1 - \rho) \cdot \frac{1}{\gamma_u}$$

$$T(\rho) = \tau \cdot \bar{T}(\rho) = \tau \left(\frac{\rho}{\gamma_l} + \frac{1 - \rho}{\gamma_u} \right)$$

This model is intentially constrained to be static in order to aid interpretation. The expectations that the agent forms based on future turnover, based on retention probabilities, encapsulates a static version of this inherently dynamic process.

Firm Objective. The firm chooses ρ to maximize:

$$\pi(\rho; S, \delta) = AL^{\alpha} C(S, \rho) - T(\rho) - wL$$

Each worker at the firm gets paid a uniform wage w, and the size of the team remains fixed at L. The first-order condition (FOC) for profit maximization is:

$$\frac{d\pi}{d\rho} = AL^{\alpha} \cdot \frac{dC}{d\rho} - \frac{dT}{d\rho} = 0$$

$$\frac{dC}{d\rho} = -\kappa\eta \cdot \text{sign}(\rho - S) \cdot |\rho - S|^{\eta-1}$$
$$\frac{dT}{d\rho} = \tau\left(\frac{1}{\gamma_l} - \frac{1}{\gamma_u}\right) = -\tau \cdot \frac{\delta}{\gamma_l\gamma_u}$$

A closed form solution for ρ^* can be arrived at as follows.
Case 1: $\rho^* > S$

$$\begin{aligned}
&\implies AL^\alpha \cdot \frac{dC}{d\rho} = \frac{dT}{d\rho} \\
AL^\alpha \cdot \kappa\eta \cdot \gamma_l\gamma_u \cdot (\rho - S)^{\eta-1} &= \tau\delta \\
AL^\alpha \cdot \kappa\eta \cdot (\delta + \gamma_u)\gamma_u \cdot (\rho - S)^{\eta-1} &= \tau\delta \\
&\implies \rho^* = S + \left(\frac{\tau\delta}{AL^\alpha \cdot \kappa\eta \cdot (\delta + \gamma_u)\gamma_u}\right)^{\frac{1}{\eta-1}}
\end{aligned}$$

Case 2: $\rho^* < S$

$$\begin{aligned}
&\implies AL^\alpha \cdot \frac{dC}{d\rho} = \frac{dT}{d\rho} \\
-|\rho - S|^{\eta-1} &= \frac{\tau\delta}{AL^\alpha \cdot \kappa\eta \cdot \gamma_l\gamma_u}
\end{aligned}$$

This case only has a solution when $\delta < 0$.
Combining the two cases together, we get the following general expression:

$$\rho^* = S + \text{sign}(\delta) \cdot \left(\frac{|\tau\delta|}{AL^\alpha \cdot \kappa\eta \cdot (\delta + \gamma_u)\gamma_u}\right)^{\frac{1}{\eta-1}}$$

Figure F.1 provides a visual representation of the marginal benefit and marginal cost of referral allocations with respect to the exogenous composition of the team. The model delivers an equilibrium where these curves intersect, wherein the firm cannot profitably allocate referrals differentially.

Comparative Statics

I analyze how the optimal referral share ρ^* responds to model primitives using the closed-form solution:

$$\rho^* = S + \text{sign}(\delta) \cdot \left(\frac{|\tau\delta|}{AL^{\alpha}\kappa\eta\gamma_U(\gamma_U + \delta)} \right)^{\frac{1}{\eta-1}}$$

When $\delta = 0$, i.e. there is no difference in the expected retention probabilities across LCs and UCs, the optimal share of referral allocation follows the (exogenously determined) share of LC workers on the team, which can be interpreted as the boundary solution $\rho^* = S$. In this section, I will focus on propositions that apply when $\delta \neq 0$, wherein the interior solution is only defined for $\rho^* > S$ if $\delta > 0$, and $\rho^* < S$ if $\delta < 0$.

Proposition 1: Effect of exogenous team diversity S Ceteris paribus, the optimal referral share ρ^* is strictly increasing in S:

$$\frac{d\rho^*}{dS} = 1 > 0$$

Proof. From the closed-form solution, ρ^* is linear in S with slope 1. □

In teams that are exogenously populated with a larger share of LCs, it is optimal for the firm to allocate a higher share of referrals to LCs. This is intuitive, since the team does not incur the cohesion cost of deviating from the inherent diversity of the team, while still capturing the retention benefits of increased representation of LCs. Figure F.2 provides a visual representation of this relationship, illustrating that the marginal benefit of referrals to LCs declines more slowly in teams that have more LCs.

Proposition 2: Effect of expected retention advantage δ The optimal referral share ρ^* is strictly increasing in δ for all values where $\delta \neq -\gamma_u$, i.e. $\gamma_l = 0$, and $\eta > 1$. That is,

$$\frac{d\rho^*}{d\delta} > 0 \quad \text{for all } \delta \in \mathbb{R},\ \gamma_l \neq 0, \eta > 1$$

In particular, for $\delta > 0$, higher LC retention raises the share of referrals going to LCs ($\rho^* > S$) above their exogenous share of the team. For $\delta < 0$, lower LC retention lowers the referral share below their exogenous share of the team ($\rho^* < S$). $\rho^* \to S$ as $\delta \to 0$, and ρ^* is continuous in δ.

Proof. We begin with the generalized closed-form:

$$\rho^* = S + \text{sign}(\delta) \cdot \left(\frac{|\tau\delta|}{AL^\alpha \kappa\eta\gamma_u(\gamma_u + \delta)} \right)^{\frac{1}{\eta-1}}$$

$$\Phi(\delta) := \left(\frac{|\tau\delta|}{AL^\alpha \kappa\eta\gamma_u(\gamma_u + \delta)} \right)^{\frac{1}{\eta-1}}$$

$$\frac{d\rho^*}{d\delta} = \text{sign}(\delta) \cdot \frac{d\Phi}{d\delta} + \Phi(\delta) \cdot \frac{d}{d\delta}\text{sign}(\delta)$$

Because $\text{sign}(\delta)$ is constant on either side of zero and discontinuous at zero, the derivative simplifies to:

$$\frac{d\rho^*}{d\delta} = \text{sign}(\delta) \cdot \frac{d\Phi}{d\delta} \quad \text{for } \delta \neq 0$$

Now differentiating $\Phi(\delta)$ for $\delta > 0$:

$$\frac{d\Phi}{d\delta} = \frac{1}{\eta - 1} \cdot \left(\frac{\tau\delta}{AL^\alpha \kappa\eta\gamma_u(\gamma_u + \delta)} \right)^{\frac{2-\eta}{\eta-1}} \cdot \left(\frac{1}{(\gamma_u + \delta)^2} \right) > 0$$

A symmetric argument applies for $\delta < 0$ with the same magnitude but opposite direction.

Hence, for all $\delta \in \mathbb{R} \setminus \{-\gamma_u\}$ and $\eta > 1$, we have:

$$\frac{d\rho^*}{d\delta} = \frac{1}{\eta - 1} \cdot \left(\frac{|\tau\delta|}{AL^\alpha \kappa\eta\gamma_u(\gamma_u + \delta)} \right)^{\frac{2-\eta}{\eta-1}} \cdot \frac{\gamma_u^2}{(\gamma_u + \delta)^2} > 0$$

□

When there are large gains from allocating referrals to LCs in (avoided) turnover and training costs, it is optimal for the firm to allocate a larger share of referrals to LCs. Theoretically, if this advantage reverses direction, it would be optimal for the firm to instead allocate a greater share to UCs. This relationship exists for cases where there are weakly increasing returns to cohesion from maintaining diversity, i.e. $\eta > 1$. Figure F.4 provides a visual representation of this relationship.

Proposition 3: Effect of cohesion sensitivity κ The optimal referral share ρ^* is decreasing in the cohesion sensitivity κ when there is a retention advantage for LCs, i.e. $\delta > 0$, and increasing in the cohesion sensitivity κ when there is a retention advantage for UCs, i.e. $\delta < 0$, when there are increasing returns in cohesion from maintaining pre-existing diversity, i.e. $\eta > 1$.

Proof. We begin with the generalized closed-form:

$$\begin{aligned}\rho^* &= S + \text{sign}(\delta) \cdot \left(\frac{|\tau\delta|}{AL^\alpha \kappa\eta\gamma_u(\gamma_u + \delta)} \right)^{\frac{1}{\eta-1}} \\ &= S + \text{sign}(\delta) \cdot \nu \cdot \kappa^{\frac{1}{1-\eta}}, \text{where } \nu \text{ is some constant} > 0 \\ \implies \frac{d\rho^*}{d\kappa} &= \text{sign}(\delta) \cdot \frac{\nu}{1-\eta} \cdot \kappa^{\frac{\eta}{1-\eta}}\end{aligned}$$

For $1-\eta < 0$, this implies that $\frac{d\rho^*}{d\kappa} < 0$ when $\text{sign}(\delta) > 0$, and $\frac{d\rho^*}{d\kappa} > 0$ when $\text{sign}(\delta) < 0$. □

κ mediates the curvature of the marginal benefit of referral allocations, resulting in a kinked design where the direction of the slope reverses symmetrically as the referral share deviates from the exogenous share of LCs on the team. As a result, if κ is low, i.e. the team is less sensitive to changes in composition, it is optimal for the firm to allocate more referrals to LCs. Figure F.3 provides a visual representation of this relationship.

Proposition 4: Effect of curvature parameter η The effect of η on ρ^* is non-monotonic, increasing for low values of η, and decreasing for large values of η, with an interior maximum.

Proof. Reframing the generalized closed-form expression:

$$\begin{aligned}z(\eta) &= \frac{\tau\delta}{AL^\alpha \kappa\eta\gamma_U(\gamma_U + \delta)} > 0 \quad \text{so that} \quad \rho^*(\eta) = S + z(\eta)^{1/(\eta-1)} \\ \frac{d\rho^*}{d\eta} &= \frac{d}{d\eta}\left(z^{1/(\eta-1)}\right) = z^{1/(\eta-1)} \cdot \frac{d}{d\eta}\left(\frac{1}{\eta-1} \cdot \log z\right) \\ &= z^{1/(\eta-1)} \cdot \frac{1}{(\eta-1)^2} \cdot \left[1 - \eta - \log(z^\eta)\right] \\ &= \left(\frac{\tau\delta}{AL^\alpha \kappa\eta\gamma_U(\gamma_U + \delta)} \right)^{\frac{1}{\eta-1}} \cdot \frac{1 - \eta - \eta \log\left(\frac{\tau\delta}{AL^\alpha \kappa\eta\gamma_U(\gamma_U+\delta)}\right)}{\eta(\eta-1)^2}\end{aligned}$$

Note that,

$$1 - \eta - \eta \log(z),$$

can be positive or negative depending on η. When η is small, $\log(z)$ is small and the expression can be positive. When η grows large, $-\eta \log z$ dominates and the derivative turns negative.

Hence, $\rho^*(\eta)$ rises at first and then falls, resulting in a unique interior maximum.

□

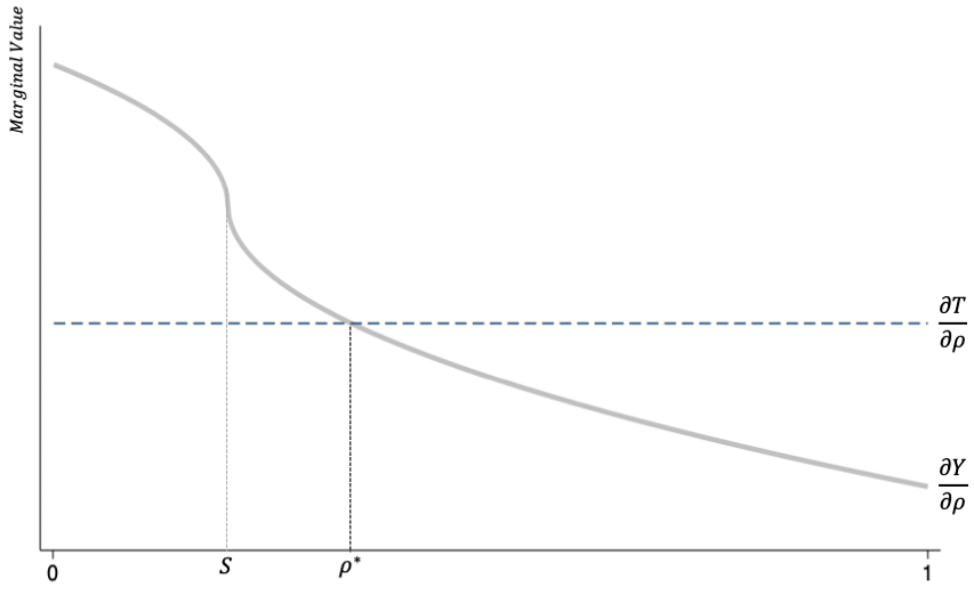
A high η implies a brittle, sensitive team culture where cohesion suffers sharply if referral shares diverge from the group's current composition. A low η suggests a more resilient or adaptable team culture where diversity imbalances are less costly in terms of cohesion.

Interpretation

The firm balances cohesion-driven productivity against retention-driven cost savings. The optimal referral strategy reflects both: referrals that are too imbalanced (relative to team composition) hurt cohesion, but referrals toward more stable workers (i.e., LCs when $\delta > 0$) reduce training costs. The equilibrium referral share ρ^* lies where marginal cohesion gain equals marginal training cost.

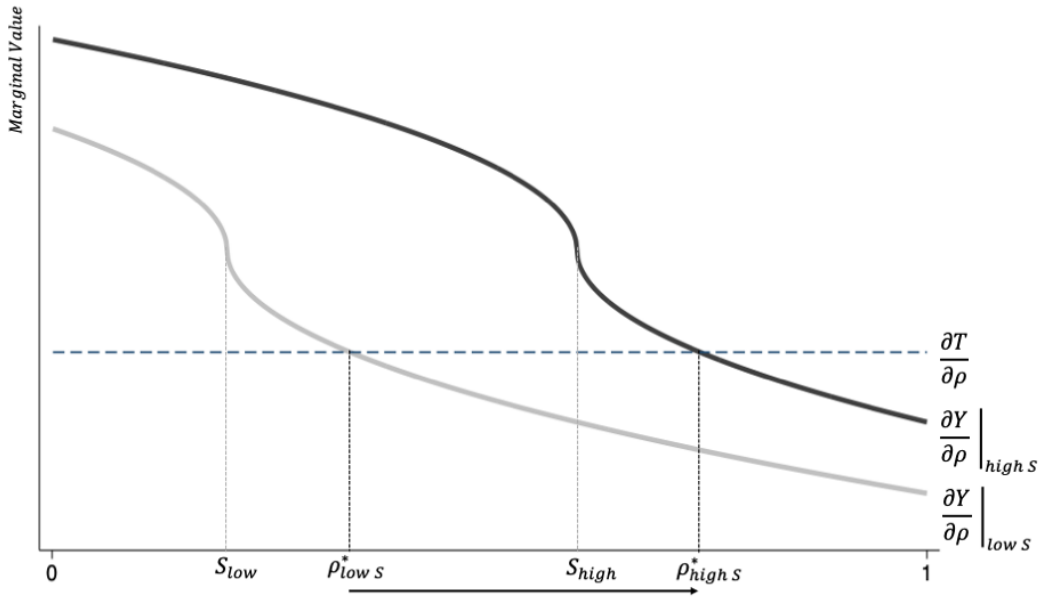
Even for a wide range of "ideal" parameter values (i.e. large δ and small κ), it is not clearly optimal to allocate all referrals to LCs, as in the experiment at the firm. This is the case for a very restricted parameter space.

Figure F.1: Static framework



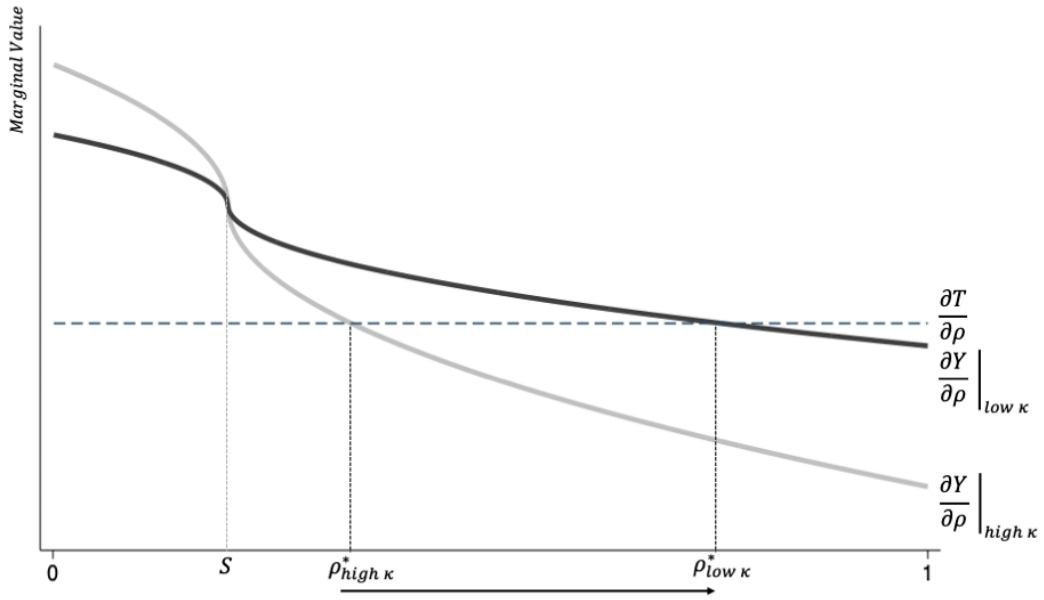
This figure represents an example of the static model framework for $\eta = 1.5$ and $S = 0.2$. The downward sloping marginal benefit curve and the horizontal marginal cost curve. They intersect at the optimal referral allocation to LCs = ρ^*.

Figure F.2: Referral allocation with respect to exogenous diversity



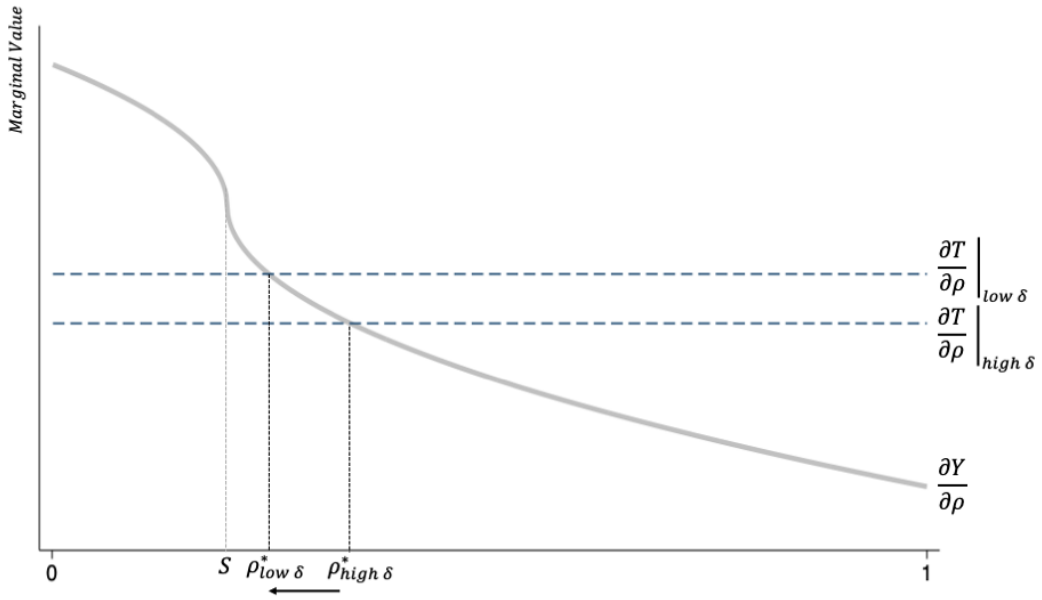
This figure represents an example of the static model framework for $\eta = 1.5$ and $S_{low} = 0.2$ and $S_{high} = 0.6$. The optimal share of LC referrals increases as the exogenous diversity of the team increases.

Figure F.3: Referral allocation with respect to cohesion sensitivity



This figure represents an example of the static model framework for $\eta = 1.5$ and $\kappa_{low} = 0.1$ and $\kappa_{high} = 0.4$. The optimal share of LC referrals increases as the sensitivity to cohesion costs decreases.

Figure F.4: Referral allocation with respect to retention advantage



This figure represents an example of the static model framework for $\eta = 1.5$ and $\delta_{low} = 0.1$ and $\delta_{high} = 0.4$. The optimal share of LC referrals decreases as the retention advantage of hiring LCs decreases.

G Bayesian learning and perceived profits

This appendix consolidates the formal dynamic model, the belief-updating machinery, and the perceived-profits simulator used for estimation.

G.1 Environment, state and flows

Time is discrete in weeks $t \in \mathbb{Z}$, with $t = 1$ the first treatment week and $t \leq 0$ pre-experiment. Let L_t^ℓ and L_t^u denote the stock of *productive* LC and UC workers. New hires enter a k_{delay}-week bench before becoming productive. Define the LC share among productive workers $s_t = L_t^\ell/(L_t^\ell + L_t^u)$ (with $0/0 := 0$). Monthly entry flows are taken from the data and distributed evenly over weeks.

The referral rule is $r_t^{\cdot} \in [0,1]$:

$$r_t^C = s_t \quad \text{(control)}; \qquad r_t^T = 1 \quad \text{(treatment)}.$$

Weekly LC and UC entries are $E_t^\ell = r_t^{\cdot} E_t$ and $E_t^u = (1 - r_t^{\cdot})E_t$. Incumbent exits follow tenure-specific hazards; let $h_{g,\tau}^{\text{mo}}$ be monthly hazards by $g \in \{\ell, u\}$ and $\tau \in \{0,1\}$ (new vs. incumbent), mapped to weekly via

$$h_{g,\tau}^{\text{wk}} \;=\; 1 - \left(1 - h_{g,\tau}^{\text{mo}}\right)^{1/4}.$$

Let $X_{g,t}^{\text{inc}} \sim \text{Binomial}\left(L_t^g,\, h_{g,1}^{\text{wk}}\right)$ and $X_{g,t}^{\text{new}} \sim \text{Binomial}\left(E_t^g,\, h_{g,0}^{\text{wk}}\right)$ denote exits among incumbents and first-week (bench) exits. Productive stocks update as

$$L_{t+1}^g \;=\; L_t^g \;+\; \text{Mature}_t^g \;-\; X_{g,t}^{\text{inc}} \;-\; X_{g,t}^{\text{new}}, \tag{9}$$

where Mature_t^g is the cohort leaving the bench after k_{delay} weeks. Weekly churn is

$$C_t \;=\; E_t + X_{\ell,t}^{\text{inc}} + X_{u,t}^{\text{inc}} + X_{\ell,t}^{\text{new}} + X_{u,t}^{\text{new}}. \tag{10}$$

G.2 Production and cohesion

In the high-frequency simulator (weekly), I parameterize effective labor as a convex aggregation of LC and UC stocks,

$$L_t^{\text{eff}} \;=\; \omega L_t^\ell + (1-\omega) L_t^u, \qquad \omega \in [0,1],$$

and use a parsimonious belief multiplier to capture the cohesion penalty:

$$y_t \;=\; A\,(L_t^{\text{eff}})^\alpha \cdot \phi(\nu_t) \;+\; \varepsilon_t, \qquad \phi(\nu_t) \equiv \frac{1}{1 + \phi_{\text{belief}}\,\nu_t} \in (0,1], \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2). \tag{11}$$

This nests the formulation $Y_t = A\, C(S_t, \rho_t) L_t^{\alpha} K_t^{\beta}$ shown in Section 5 by holding K_t fixed at the weekly horizon and mapping $C(\cdot)$ to $\phi(\nu_t)$.

G.3 Beliefs about cohesion costs

Structural Bayesian updating. Supervisors enter with Gaussian priors

$$\kappa \sim \mathcal{N}(\nu_{\kappa,t},\, \sigma^2_{\kappa,t}), \tag{12}$$

and observe a per-referral cohesion signal $s_t = (Y_t - \bar{Y}_0)/N_t$ when $N_t > 0$ LC hires occur. Under (6), $s_t \mid \kappa \sim \mathcal{N}(\kappa, \sigma^2_\varepsilon/N_t)$. Once a minimal stock of LC referrals $\sum_{\tau \le t} N_\tau \ge k^*$ is accumulated, beliefs update via Normal–Normal conjugacy:

$$\sigma^2_{\kappa,t+1} = \left(\sigma^{-2}_{\kappa,t} + N_t/\sigma^2_\varepsilon\right)^{-1}, \tag{13}$$

$$\nu_{\kappa,t+1} = \sigma^2_{\kappa,t+1}\left(\sigma^{-2}_{\kappa,t}\,\nu_{\kappa,t} + N_t s_t/\sigma^2_\varepsilon\right), \tag{14}$$

and otherwise remain at $(\nu_{\kappa,t}, \sigma^2_{\kappa,t})$.

Reduced-form belief path for weekly profits. For estimation and to align timing with the reduced-form moments, I use a smooth parametric path that approximates the Bayesian learning dynamics:

$$\nu_t \;=\; \kappa_{\text{true}} \;+\; \left(\kappa_{\text{prior}} - \kappa_{\text{true}}\right) \exp\!\big[-\eta\,\big(t - t_{\text{learn}}\big)_+\big], \qquad (x)_+ \equiv \max\{x, 0\}, \tag{15}$$

with speed $\eta > 0$ and delay $t_{\text{learn}} \ge 0$. A stock threshold k^* can accelerate the drift once enough LC referrals have accumulated. The mapping from ν_t to output is via $\phi(\nu_t)$ in (11).

G.4 Profits and frictions

Perceived weekly profits are

$$\pi_t \;=\; y_t \;-\; \mu C_t \;-\; f \;-\; \mathbf{1}\{t \ge 1\}\, c_0 e^{-t/\tau_{\text{impl}}} \;-\; \lambda\,|\rho_t - s_t|, \tag{16}$$

where μ weighs turnover/training costs, f is a fixed operating cost, c_0 and τ_{impl} capture a fading implementation friction specific to the treatment arm, and $\lambda|\rho_t - s_t|$ penalizes short-run reorganization when the referral rule departs from the status quo composition. This profit mapping is the modified weekly version of the common overarching framework discussed in the main paper in equation (7).

G.5 Estimation by simulated method of moments

Let $\theta = \left(A, \alpha, \mu, f, \kappa_{\text{true}}, \kappa_{\text{prior}}, \phi_{\text{belief}}, \eta, t_{\text{learn}}, k^*, c_0, \tau_{\text{impl}}, k_{\text{delay}}, \omega, \sigma_\varepsilon\right)$. Given θ, the simulator (i) takes monthly entries from the data and weekly hazards $h^{\text{wk}}_{g,\tau}$, (ii) updates stocks and beliefs week by week using (9)–(15), and (iii) produces arm-by-time series for output, turnover and referrals. I form simulated moments to mirror:

(M1) Event-time means of standardized weekly output (treatment and control).

(M2) Arm-by-month means of entries, exits and churn.

(M3) Arm-by-month shares of LC referrals and the implied LC composition.

Parameters minimize $Q(\theta) = [m_{\text{data}} - m^b(\theta)]^\top W [m_{\text{data}} - m^b(\theta)]$ with a diagonal weight matrix W equal to the inverse of the estimated variance of each moment (team-level variance / number of teams). Uncertainty bands in Figure 4 are Monte Carlo envelopes over binomial exit and Gaussian output shocks; the mean crossing time $T^* = \min\{t \geq 1 : \bar{d}_t > 0\}$ is about 6.2 weeks with a 10–90 percentile interval [5, 14].

Equation (15) produces the belief paths underlying Figure 3, consistent with the Normal–Normal learning logic in (13)–(14) when treatment forces early exposure to LC referrals. Combining (11) and (16) with observed entry flows and hazards yields the perceived profit trajectories in Figure 4: a short initial dip followed by recovery as beliefs improve and LC cohorts mature into productive work.