The London School of Economics and Political Science

Essays on the allocation, coordination, and selection of workers

Virginia Magda Luisa Minni

A thesis submitted to the Department of Economics
for the degree of Doctor of Philosophy

London, March 2023

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorization does not, to the best of my belief, infringe on the rights of any third party.

I declare that my thesis consists of approximately 70,000 words.

Statement of conjoint work

I confirm that Chapter 2 was jointly co-authored with Laura Boudreau, Rocco Macchiavello, and Mari Tanaka, and I contributed 30% of this work. I confirm that Chapter 3 was jointly co-authored with Nava Ashraf, Oriana Bandiera, and Victor Quintas-Martinez, and I contributed 30% of this work.

Abstract

This thesis studies the determinants and consequences of workers' allocation, coordination, and selection within organizations in countries at different levels of economic development.

The first chapter provides evidence of the critical role of managers in matching workers to jobs within the firm using the universe of personnel records on 200k employees over ten years of a multinational firm. Leveraging exogenous variation induced by the rotation of managers across teams, I find that successful managers cause workers to reallocate within the firm through lateral and vertical transfers. This leads to large and persistent gains in workers' career progression and productivity. The results imply that the visible hands of managers match workers' specific skills to specialized jobs, leading to an improvement in the productivity of existing workers that outlasts the managers' time at the firm.

The second chapter continues the study of leadership in a very different context: Myanmar's labor movement. We conduct multiple field experiments by collaborating with a confederation of labor unions as it mobilizes garment workers in the run-up to a national minimum wage negotiation. First, we document that union leaders differ from the other workers along several traits that psychologists and sociologists have associated with the ability to influence collective outcomes. Second, by randomly embedding leaders in group discussions, we find that they help coordinate workers' views to build consensus around the unions' preferred minimum wage levels. Third, by conducting a mobilization experiment that features collective action problems, we show that leaders play a coordinating role also for workers' actions.

The third chapter starts with the fact that women's labor force participation varies widely across countries at every level of development. We ask how this affects gender diversity among employees, gender gaps, and firm productivity using five years of personnel records on over 100K employees of a multinational firm combined with the female to male labor force participation rate in the 101 countries where the firm operates. Structural estimates show that in a counterfactual world with no gender-specific barriers to labor force participation, firm productivity would be 32% higher for the same level of employment and the same wage bill. The findings suggest that selection is a powerful lens to understand the link between diversity and productivity.

Acknowledgements

My supervisors have provided invaluable guidance throughout my PhD and greatly contributed to my development as a researcher. I am sure my research path would have been much different had I not had such an extraordinary committee. Oriana Bandiera took a chance on me in my first year of graduate school when I showed up in her office, with very limited research experience, and with eyes beaming when thinking about doing empirical research in collaboration with organizations. Her unconventional brilliance and contagious creativity inspire me every day. John Van Reenen always pushed me toward excellence. He listened to my countless research ideas and presentations; each time, he would promptly come up with practical advice on how to proceed further which reflects his passion for everything economics and his extraordinary knowledge of methods and papers outside his core research agenda. I am deeply thankful for his unfailing help in thinking through many challenging, and sometimes tedious, economic problems with me. Nava Ashraf empowered me to rise to the occasion when collaborating with organizations and helped me find the right words and the confidence to introduce myself to the academic world. She exposed me to countless stimuli from economics, the social sciences, and the arts that enhanced my life. She is an exceptional example of impeccable intuition and authentic human connection. Robin Burgess became my official supervisor the day he told me that I had the potential to succeed in academia but that I had to dare to face the unknown and overcome the fear and the sense of vertigo. His skill to grasp the essence of projects and come up with a paper title in the blink of an eye is nothing short of magic.

I have also benefited from the advice and input of many others. I thank the people in the organizations that collaborated with me: without their support and curiosity none of this research would have been possible. I feel enriched by all the conversations I had with Tim Besley, Laura Boudreau, Gharad Bryan, Maitreesh Ghatak, Gilat Levy, Rocco Macchiavello, Isabela Manelici, Alan Manning, Guy Michaels, Steve Pischke, Nina Roussille, Yona Rubinstein, and Catherine Thomas. I am very grateful to the Economic and Social Research Council (ESRC) for my PhD scholarship and to STICERD for financial support during my PhD. For the second chapter, I gratefully acknowledge financial support from the IGC and the Institute of Economic

Research at Hitotsubashi University.

I have shared this ride with friends and colleagues that have provided essential support and inspiration. Gaia Dossi has been the best friend I could ever hope for, we shared so many meaningful conversations about research and life. I would also like to especially thank Alix Bonargent, Amanda Dahlstrand, Alexia Delfino, Edoardo Leonardi, Akash Raja, Veronica Salazar-Restrepo, Victor Quintas-Martinez, Martina Zanella, and Celine Zipfel not only for insightful comments on my work but also for contributing to making my PhD years memorable.

I am deeply grateful to Roberta for her presence and for making everything worthy of attention.

My parents and grandparents have always been forgiving of my stubbornness. They taught me that I could do whatever job I wanted, as long as I stayed true to myself, treated others with respect, and always remembered that life was larger than myself. They are incredible human beings, and I am so lucky to have them around. The quiet and wise aura of Sofi comes from another planet, and it is the only thing I really need most of the time. Bianca's self-reliant drive and energy invariably fill me with joy.

My partner Stefano has been by my side for 12 years. He gives me unconditional love and support and at the same time urges me to evolve like no one else can. We have grown so much together.

This thesis marks the end of a long, intense, and wild journey. And also the start of a new one.

This one's for me.

Contents

List of Figures

List of Tables

Chapter 1

Making the Invisible Hand Visible: Managers and the Allocation of Workers to Jobs

"[M]odern business enterprise took the place of market mechanisms in coordinating the activities of the economy and allocating its resources. In many sectors of the economy, the visible hand of management replaced what Adam Smith referred to as the invisible hand of market forces."

— Chandler, A.D., 1977. *The Visible Hand: The Managerial Revolution in American Business.*

1.1 Introduction

Economics studies how to allocate scarce resources. Traditionally, labor economics focused on the labor market, rather than looking inside the "black box" of firms, within which most workers are allocated to jobs.[1] In firms, managers take the place of the price mechanism in directing the allocation of resources (Coase, 1937). In particular, they shape the allocation of workers to jobs through *internal labor markets* (Doeringer and Piore, 1971). Understanding the managers' role in the allocation of workers to jobs is key to understanding why differences in management across and within firms explain an important share of the persistent differences in productivity (Gibbons and Henderson, 2012).

The idea that there are gains from the division of labor with people specializing their efforts across tasks is an old one and among the cornerstones of economics (Smith, 1937). Yet, the matching of workers to jobs as a way to reach an organization's objectives has received little attention. Managers, acting as gatekeepers in internal labor markets (the *bosses*), can play

[1]The share of workers employed by firms is 54% globally (World Bank, 2019).

an essential role in facilitating the discovery of workers' unique skills and hence their effective utilization through job allocation.

This paper documents how managerial skill shapes workers' allocation to jobs and future career outcomes and whether this ultimately determines firm productivity. I consider a setting that allows the study of workers' career trajectories both horizontally - through lateral moves - and vertically - through a job ladder. This is the internal labor market of a large multinational firm (MNE).

Studying the role of managers within internal labor markets requires tackling three steps. The first is access to "insider" firm data, which also combines cross-sectional granularity with a sufficiently long time dimension. Second, estimating the added value of managers has proven challenging as measures that identify good managers independently of workers' outcomes are hard to come by. Third, to analyze the impact of managers on workers, one needs to pin down the manager's contribution to the worker's outcomes, which necessitates plausibly exogenous assignment of managers to workers.

With respect to the data, I bring together a rich collection of high-granularity administrative records from a multi-billion euro multinational firm. The data reveal the organization's inner workings over several years and cover the universe of managers and workers in the MNE: more than 200,000 workers and 30,000 managers over the span of 10 years in 100 countries.

To address the first identification step, I advance a new method to identify successful managers based on managers' own promotion speed, as a revealed preference measure of the firm. I refer to them as "high-flyers" to capture those that climb the corporate ladder faster. Specifically, I consider the earliest age a worker is promoted to manager and define a binary measure to classify managers as high-flyers and low-flyers. This results in 29% of managers being singled out as high-flyers.[2]

I tackle the second identification step by leveraging a *natural experiment* created by managers' lateral rotations across teams that are outside of the control of the worker and adopt an event study strategy. These rotations are part of the requirement for the managers' career progression and anecdotal evidence and empirical tests indicate that they are orthogonal to workers' characteristics.[3] This type of rotation policy is also not peculiar to this firm but rather

[2] I show that the high-flyer status is significantly positively correlated with other measures of ex-post performance such as managers' own performance appraisals as well as workers' upward feedback on the managers' leadership.

[3] I carry out a series of empirical tests to confirm that the team assignment of a manager's internal rotation is orthogonal to workers' characteristics. I show that the type of transition faced by a team (e.g., from a low- to a high-flyer manager) is uncorrelated with the observable characteristic of the team, as well as with the characteristics of the incoming and outgoing managers. Most importantly, the career progression of workers undergoing different manager transitions follows parallel trends leading up to each type of manager transition.

common managerial practice among large firms.[4] I conduct an event-study analysis exploiting the workers' first manager rotation and comparing different types of transitions. For example, consider two teams each managed by a low-flyer manager. One of these teams then transitions from the low-flyer manager to a high-flyer manager, and the other team transitions from the low-flyer manager to a different low-flyer manager. As both teams are affected by a manager transition, this design nets out the effect of the transition. Hence, the results can be summarized in the effects of (i) *gaining a good manager*, i.e. switching from a low- to a high-flyer manager, and (ii) *losing a good manager*, i.e. switching from a high to a low-flyer manager, relative to switching manager but without changing manager type. I can compare the outcomes of the employees each month leading up to the manager transition date and each month after the transition.[5]

I show that good managers achieve a more productive workforce by creating better matches between the present labor pool and specialized jobs in the firm. In so doing, they have a long-lasting impact on workers' trajectories that outlives their time overseeing the worker.[6] My findings suggest that considerable gains in worker performance stem from efficiently allocating *existing* workers to jobs and that managers' role is crucial in creating more productive worker-job matches, all potentially at little additional cost for the organization.[7] As the managers' influences propagate inside the organization through their subordinates' careers, I demonstrate that they significantly impact firm-level productivity, thus linking individual-level effects to the productivity of an entire establishment.

First, gaining a good manager causes significant worker reallocation to different jobs inside the firm, through lateral transfers (30% higher) and also vertical transfers (40% higher). Examples of lateral moves are transfers from customer service to logistics; from merchandising to sales; or from product development to quality. Moreover, I isolate task-distant transfers as those that represent a major horizontal change in tasks to be fulfilled and find that they increase by 20% (for e.g., moving from human resources to marketing, or from R&D to supply chain management). I find no systematic pattern among the moves, they are scattered throughout the organization. In terms of dynamics, the transfers gradually increase until five years after the manager transition when they level off at a sustained higher level at least until seven years after.[8] The results on

[4]Other studies have exploited similar rotations in different organizations for their identification strategy (see e.g., Cullen and Perez-Truglia (2019); Haegele (2021)). Besides firms in the private sector, rotation policies are also used in large public organizations such as the World Bank and the United Nations.

[5]I keep following the workers even if they change managers again, irrespective of whether the worker remains or not with the manager of the first transition.

[6]Having panel data over several years is essential to be able to evaluate the returns of a worker-job match as they may not manifest immediately.

[7]Matching can be considered a resource-neutral policy when contrasted to the more resource-intense alternatives such as hiring, firing, and training.

[8]The time window is determined by the length of the panel data.

the lateral transfers cannot be reconciled with high-flyer managers mainly teaching workers on how to become more productive on-the-job as that would lead to the opposite prediction on workers' lateral moves.[9]

Second, gaining a good manager also results in an improvement in worker performance and long-run career progression. Seven years after the manager transition, the number of salary grade increases is 0.25 points higher, corresponding to a 30% higher salary. Combining the results on the lateral reallocation with those on pay progression suggests that high-flyer managers facilitate the discovery of workers' aptitudes and spur workers to a higher rate of job changes, which results in workers finding positions that are better matched to their skills. A mediation analysis reveals that 62% of the higher salary grade increases are explained by lateral job changes. This is likely an underestimate of the managers' matching channel since it excludes vertical transfers (as by definition they involve a salary raise) and it also does not consider the gains due to a worker remaining in the current job because it is a good match.

Third, using objective productivity data from sales bonuses on a sub-sample, I show that good managers boost worker performance, rather than inflating pay for the same performance. I find that workers' sales performance increases by 27% up to 4 years after gaining a high-flyer manager.[10] Additional empirical checks that compare the productivity gains among job moves initiated by a high-flyer with those from job moves initiated by a low-flyer indicate that the performance gains cannot be explained by a treatment effect of transfers by themselves, but rather by good managers causing more productive transfers (i.e. choosing the *right* transfer for the *right* worker in terms of the worker's skill set).

These effects are *asymmetric*. Gaining a good manager has positive effects while losing one has no corresponding negative effects. This indicates that there are long-term benefits of a one-time exposure to a good manager: the gains from a high-flyer manager persist even after a downgrade in a manager's capacity. The asymmetric effects together with the persistence of the results help rule out alternative contemporaneous channels of managers such as monitoring or motivation and support the interpretation of the matching channel as the gains of a good worker-job match do not rely on the co-presence of a good manager. In terms of organizational design, the asymmetries in the results also indicate that it suffices to expose each worker to a high-flyer once as a low-flyer manager cannot spoil away the benefits of a good match created

[9] I show this formally with a conceptual framework that captures task-specific human capital and learning about innate talents. I allow good managers to increase both the learning around task talent (matching channel) and the speed of job-specific learning by doing (teaching channel). I show that the two channels have opposite predictions on job transfers and that the data is consistent with the matching channel being the main driver behind the productivity results.

[10] I have objective sales bonus data for the entire field sales population in India over 2018-2021. The corresponding increase in salary is 8% in the same sales sub-sample and salary is 18% in the full sample 4 years after the manager transition.

by a high-flyer manager.

Additional tests allow me to rule out two additional alternative channels. First, I do not find that the workers' lateral and vertical moves occur within the managers' networks of previous colleagues, thus excluding explanations related to social connections within the MNE.[11] Second, the findings on worker performance cannot be explained by high-flyer managers engaging in worker selection out of the firm (Fenizia, 2022). I observe no impact on exit from the firm, and this is not disguised by heterogeneous effects on exit by baseline worker performance: there is no impact on exit for either the high or low performers at baseline. Hence, the higher rate of internal transfers points to high-flyers finding suitable re-deployments inside the firm.

I conclude by showing that the good managers' effects are increasing overall profits at the establishment level. I integrate the worker-level records with establishment-level productivity data (output per worker) and cost data (costs per ton), connecting the paths of individual workers to the overall productivity of the establishment. Although this piece of evidence is correlational in nature, it provides further evidence of a positive link between the career trajectories of individual workers and productivity at the site level. I estimate that the semi-elasticity of output per worker to workers' past exposure to high-flyer managers is 2.03, that is increasing the exposure to high-flyers by 10 percentage points is associated with an increase in output per worker by 20%. The same semi-elasticity is -1.4 for costs per ton. Taking the price level as given and combining together these two results, the analysis suggests that the high-flyers' effects are increasing profits.[12]

A major question in labor economics is how workers match to jobs and how that determines wages and their evolution over time. Extensive research on labor markets has studied job mobility *between* firms (e.g., Jovanovic (1979); Rosen (1986); Moscarini (2005); Acemoglu and Autor (2011); Bagger et al. (2014); Chade et al. (2017); Card et al. (2018); Lise and Postel-Vinay (2020)). Yet, wage growth and job mobility also happen *within* firms as examined by a literature on internal labor markets, largely theoretical and descriptive (Waldman (1984); Baker et al. (1994a); Baker et al. (1994b); Baker and Holmstrom (1995); Gibbons and Waldman (1999); Kahn and Lange (2014); Pastorino (2019)). This is the first paper to study the role of managers in the allocation of workers to jobs within internal labor markets and to show that manager quality is the crucial ingredient needed to create more productive matches between

[11] I define a socially connected move based on whether the manager has ever worked (i) with the new manager the worker moves to and/or (ii) in the same sub-function and/or office as the job the worker moves to. I find no differential impact of gaining a high-flyer manager on connected moves, whether these are lateral or vertical moves.

[12] Regarding the constant prices assumption, I am controlling for country, product category, and time fixed effects in the regression. Moreover, pricing is the responsibility of marketing teams and production managers do not set prices.

workers and jobs.

My findings also advance our understanding of the impact of individual managers on firm and worker outcomes (Bertrand and Schoar (2003); Bandiera et al. (2007); Lazear et al. (2015); Bandiera et al. (2020); Frederiksen et al. (2020); Hoffman and Tadelis (2021); Metcalfe et al. (2022); Adhvaryu et al. (2022a); Adhvaryu et al. (2022b)). I contribute to this growing strand of research by uncovering the matching of workers to jobs as an important mechanism that determines managers' long-run impacts on workers' careers and overall firm productivity. In so doing, I also bring forth new evidence on the micro-level processes that link individual managers at lower levels of the firm hierarchy to firm-level outcomes. In terms of management practices, this study puts the emphasis on managerial policies governing the allocation of workers to jobs within firms, which have been overlooked by previous research.[13]

More broadly, by providing micro-level evidence on the role of managers in the efficient assignment of workers to jobs, this study speaks to the research on the misallocation of productive inputs and growth: (i) on the mismatch between workers and jobs and its consequences for workers' careers and aggregate output (Hsieh et al. (2019); Guvenen et al. (2020)) (ii) on the misallocation of productive resources across firms in the economy and the role that the reallocation of factors of production can play in driving productivity growth (Bhagat et al. (1990); Bartelsman et al. (2009); Hsieh and Klenow (2009); Foster et al. (2001); Davis et al. (2014)).

The rest of the paper proceeds as follows. Section 2.2 describes the institutional background and Section 1.3 the data. Section 1.4 introduces the research design centered around manager rotations and discusses its validity. Section 1.5 presents the results, and Section 1.6 discusses additional evidence corroborating the matching channel and ruling out other possible channels. I develop a conceptual framework in Section 1.7 that provides a stylized theory of the role of managers in internal labor markets and helps interpret the empirical results. Section 1.8 concludes with policy implications and other issues for further research.

[13] The managerial practices analyzed by previous literature focus on workers' incentives via pay for performance, promotions, and monitoring (Bloom and Van Reenen, 2011). The tools of monetary and career incentives have also been widely examined theoretically and empirically by a prominent strand of research in organizational economics (Holmström (1979); Lazear and Rosen (1981); Lazear (2000); Bandiera et al. (2007); Bandiera et al. (2013); Bertrand et al. (2020)).

1.2 Institutional context

1.2.1 Firm overview

I collaborate with a private consumer goods multinational with offices in more than 100 countries worldwide. This firm is one of the largest in the world and is headquartered in Europe. It has a workforce of about 120,000 workers each year, of which approximately 60,000 are white collars, and its turnover in 2020 was over €50 billion. I collect novel data on the full population of white-collar and management employees and construct a panel dataset that links workers to their managers and tracks workers' career progression inside the firm.

The company is organized into a hierarchy of work-levels (WL) that goes from WL1 to WL6 (C-Suite) (see Appendix Figure 1.47 for a graphical visualization of the hierarchy). Employees with a work-level above one are considered performing managerial roles (WL2+). Moreover, within each work-level, there is a further vertical differentiation through salary grades (there are 12 salary grades in total). Salary approximately increases by 20%-30% at each salary grade increase. A salary grade increase entails a permanent change in salary but not a major change in job responsibilities while a work-level promotion would also entail a considerable change in job responsibilities (usually less execution and more strategy and planning). The firm has the same organizational structure across all countries, functions, and over the time of the sample. Appendix Figures 1.54-1.56 show that average tenure, age, and work-level shares have remained very stable over the years of the panel.

Table 1.1 describes my sample, which consists of the universe of white-collar workers from January 2011 to December 2021. This results in 205,432 distinct regular full-time workers[14] in 118 countries (8,618,267 worker-month observations). Supervisors (i.e. those that supervise at least one worker) comprise 21% of the sample, although 15% are in managerial roles[15] (i.e. have a work-level above one).

Table 2.1 presents summary statistics for the main variables. Women represent 43% of employees in the sample, 39% of workers are aged between 30-39 and the large majority of workers are in work-level 1, WL1 (81%). The workers have homogeneous levels of human capital as applications require a college degree, and most employees have degrees in either economics and business administration (49%) or STEM (30%). Tenures at the firm are long, with an average of 8.5 years, highlighting the importance of internal career progression for employees' long-term income. Teams (i.e. a group of workers reporting to the same supervisor) are small with an average of 5 workers per team, although team size increases over a manager's seniority

[14] 97% of employees work full time.

[15] Exactly the same as the global average share of managers among the white-collar workforce reported by ILO in 2019.

with top managers overseeing an average of 8 workers.

Because I am interested in career progression to higher-level positions, I focus on white-collar employees. Blue-collar workers have very limited career progression opportunities as well as horizontal job differentiation (87% of blue-collar workers are machine operators). Moreover, the organization of work in factories is very different from offices; blue-collars are supervised by white-collar front-line workers (denoted as first-line managers) instead of employees in actual managerial positions and teams that can be as large as 80 workers.

The workers and workplace practices at the firm are comparable to those of other large European manufacturing firms. On average, in firms with more than 250 workers, the gender share of the workforce is 40%, age is 41 years and tenure is 9.8 years. Moreover, the typical large firm would have at least 4 hierarchical levels.[16] As of June 2022, the most common job in the United Kingdom was a Product Developer in the R&D function with an average annual salary of GBP 39,190, very much in line with Glassdoor's average salary of GBP 39,313.[17]

1.2.2 The role of line managers

Line managers are responsible for setting team priorities, coaching, and giving feedback to workers. They can significantly influence job design through the assignment to projects inside and outside the team. Crucially, managers' input is key for promotion and transfer decisions (in line with other organizations, Frederiksen et al. (2020), Haegele (2021)). Managers have an explicit incentive to "develop and magnify the power of people". Their periodic evaluation is structured around seven "standards of leadership" and one of these is to be a "talent catalyst" who "coaches individuals and teams to realize their full potential".

The firm uses 360-degree evaluations: a line manager receives written evaluations from both superiors and subordinates on each of the indicators and his own manager reviews these to decide on a single (numerical) performance rating each year.[18] Line managers formally review their subordinates' work every quarter, where they also identify priority skills and development areas for each worker but the overall performance rating is annual. They are also encouraged to have weekly 1-1 meetings with each worker to re-assess priority and check status (see Appendix Figure 1.49 for an excerpt of the firm HR guidelines to managers).[19] In 2020, employees reported in the annual global pulse survey at the firm that their line manager was among the top three areas of

[16] European-wide statistics are taken from the European Company Surveys (van Houten and Russo, 2020).

[17] Glassdoor's page for Product Developer in the United Kingdom.

[18] The written text of these evaluations did not pass the confidentiality criterion for the data to be shared for this research as it was deemed that they could not be cleaned so to preserve anonymity. Only the numerical performance ratings were shared, which are then used to determine the annual bonus.

[19] Qualitative evidence from focus groups of workers at the firm indicates that frequent 1-1 meetings with the line manager tend to go hand-in-hand with good managers.

importance to them, further underscoring the relevance of this relationship in the workplace.

These firm policies are in line with managers' job responsibilities among white-collar employees in other companies (Clifton and Harter, 2019). In these higher-skilled, knowledge-based jobs, production is often complex and multi-faceted and firms care about both current performance and future performance, i.e. workers' "potential" and career paths (Benson et al., 2021).[20]

1.3 Data

The main variables in the analysis are obtained from the personnel records of the organization, which provide monthly snapshots of the workers worldwide. I assemble rich panel data by combining the global HR records with the organizational chart, the payroll and performance data, and the annual surveys. Appendix Figure 1.46 illustrates the various data sources and the time period for which they are available. Table 1.3 presents summary statistics for the outcomes.

1.3.1 Personnel records

The global personnel records keep track of demographic variables of interest (age, gender, tenure, education), and give a monthly snapshot of the workers' hierarchy levels (from which promotions can be constructed as described in sub-section 1.2.1), functions, and job titles. It is also recorded if a worker has been made redundant (involuntary exit) or if she has found alternative employment (voluntary exit).

In terms of the types of jobs, there are 14 functions in the MNE, with the biggest six being Sales, HR, R&D, Supply Chain, Finance, Marketing. Within each function, there are multiple sub-functions (for e.g. in the finance function one can be working in the tax sub-function or in the M&A sub-function). Typically, a sub-function would have roles spanning from work-level 1 to work-level 4, so workers do not have to change sub-function to move up the job ladder as it is possible to advance vertically within a given sub-function.[21] I also observe the job titles detailing a worker's exact job within the sub-function. There are almost 1000 horizontally differentiated job titles within the firm (heterogeneity in terms of the tasks associated with jobs at the same level of the job ladder), and, on average, there are two distinct job titles in a team.[22]

[20]The study population is knowledge-based workers as opposed to lower-skilled workers, who have been the subjects of most of the empirical personnel papers.

[21]The median size of a sub-function is 241 workers, the 10th percentile is 16 workers and the 90th percentile is 2112 workers.

[22]These are some examples of job titles: Logistic Specialist; Supply Planning Admin; HR Recruiting Specialist; Occupational Health Admin; Field Sales Specialist; Vice President Brand Development.

1.3.2 Organizational chart

The organizational chart indicates the manager each individual worker reports to, where workers reporting to the same manager belong to the same team. Because these data capture team assignments over many years, I am also able to construct measures of managers' formal ties to other units at the firm by measuring whether they have previously worked with anyone in that unit.

1.3.3 Performance data and objective productivity

I supplement this data with payroll data, capturing employees' earnings, and bonus payments.[23] Pay, which is available from 2016 onwards, captures differences in performance across workers and there is considerable variation in pay within a given job in a specific office-month pair, where the median standard variation in pay is around €6,000 (for the whole distribution see Appendix Figure 1.51). Practically, there are three ways in which workers in the same job title can earn a different pay: the salary grade[24], the salary band[25] and the annual bonus (variable pay, which is on average 10% of fixed pay for entry-level workers). Appendix Figure 1.52 shows that lateral moves are common in every sub-function and Appendix Figure 1.53 shows that this is also true for salary grade increases.

In addition, I collect information from the firm's talent management system which includes worker evaluations, such as performance ratings, set annually by the manager. Salary increases and promotions are the main metrics to assess performance within the firm. The manager is the main decision-maker after taking into account the views of all the colleagues that have interacted with the worker (360-degree reviews). The decision process is designed to be as fair as possible and to limit manager bias, the manager has to justify any salary increase, transfer, or promotion decision against a set of objective criteria to the rest of her colleagues in specific talent forums dedicated to this discussion. The performance assessment is done in the same way in every function and branch so that comparisons can be made between workers in different jobs and offices.

I complement the performance data with two independent sources of objective productivity data. The first is objective sales bonus data at the worker level for the full Indian sales population from January 2018 until December 2021 (around 2500 employees). India is the largest country in terms of overall employee size and the sales bonus is based on reaching targets each month

[23] Salary is measured in euros in all countries.

[24] Appendix Figure 1.61 shows the positive relationship between the number of salary grade increases and pay in logs.

[25] Within each salary grade, there is a salary band that goes from 80% to 120% of target pay determined via market benchmark data.

set by the country head office. Some examples of sales targets include: growth of sales, product placement, on-shelf availability, additional exhibitions, the number of orders vs total visits each month.[26] The second is operational data at the establishment level: first, output per worker (tons per FTE or Full-Time Equivalent), a common metric of productivity in manufacturing firms, and second, costs per unit of output (operational costs per ton).[27] Both of these measures are at the establishment-year level and the company shared all data available for every factory globally (around 150 sites) over 2019-2021. Because of changing reporting requirements, the costs per ton data could only be shared for the main product category (there are three product categories in total).

1.3.4 Digital platforms: career mobility and flexible projects

In 2018, the firm introduced two platforms aimed at fostering an internal talent marketplace.

The first is denoted as a career mobility platform and works as a talent tool combining learning and development, skill analytics, and career mobility. Workers can use the platform to do workshops, search for internal jobs and share learning/job opportunities. Appendix Figure 1.48 shows a mockup that replicates a hypothetical user's profile page. The data available tracks the workers' activities in the company such as the number of completed courses, number of posted skills, and number of items shared with colleagues. I use this data to infer workers' engagement in career development activities.

The second platform is a tool that enables workers to apply for short-term projects inside the company, denoted as flexible projects. These projects can vary in duration but typically range between one to six months and entail one or two days per week of work on the flexible opportunity. The rationale underlying this initiative is rooted in two objectives: to allow workers to engage in small projects to experiment with different jobs, expand and test their skills, as well as to fill new positions in real-time in response to quickly changing market needs.

1.3.5 Global employee surveys

I conduct additional analysis using individual responses to four global annual surveys that the company ran in 2017-2021. Each September, all workers are invited to the survey; the response rate is around 60%.[28] The survey is designed to measure the "pulse" of the workers across the globe, gathering data on how the organization is perceived by the workers themselves and on

[26] Appendix Figure 1.59 shows that there is a positive relationship between current productivity and future salary grade increase. In particular, a 1s.d. increase in sales bonus over a year increases the probability of an increase in the salary grade in the next year by 80%.

[27] The operational costs are predominantly made up by labor and energy costs and they do not include the cost of raw materials.

[28] In 2017 and 2018, the survey was only sent to a random sample of employees.

their job satisfaction and well-being (questions are on a 5-point Likert scale). Respondents are broadly similar to non-respondents in terms of demographics; they generally tend to be slightly older, higher up in the hierarchy, and are marginally more likely to have a high-flyer manager (Appendix Table 1.25).

1.4 Empirical strategy

My analysis revolves around the causal effects of high-flyer managers on the subsequent career progressions of their workers. For example, I want to measure whether workers fare better after transitioning from a low- to a high-flyer manager (e.g. have higher wage growth). To estimate these manager effects I would ideally randomize employees to their managers. As this type of experiment is not feasible, I instead exploit naturally occurring exogenous rotations in manager assignments within the organization (natural experiment). I first describe how I identify high-flyers and the manager transitions, and then specify the formal econometric framework for the event-study analysis.

1.4.1 High-flyers

I construct a new proxy for good managers based on managers' own speed of promotion. It is a measure of the managers' personal success in the organization and it is not directly based on the outcomes of their workers.[29] In particular, I define high-flyer managers as those that achieve their maximum observed work-level at a relatively younger age (time-invariant). I only look at work-level 2 managers since the focus of the paper is on middle managers, the major part of the managerial workforce in large firms (see Figure 1.1 for the distribution of work-levels at different tenure years). Because of data confidentiality, I only observe 10-year age groups. Hence, the high-flyer measure is defined as achieving work-level 2 below the age of 30. Overall, 29% of managers are high-flyers (Figure 1.2 shows the full distribution over age group). The share of high-flyers is broadly constant across functions, countries, and years.

The intuition behind this measure is that the speed at which one progresses the corporate ladder is symptomatic of leadership potential and it reflects the extent to which the firm values the manager's work.[30] I validate this intuition empirically by showing that the high-flyer status is significantly positively correlated with other measures of performance. First, Figure 1.3

[29]Previous studies have based their measure of manager quality directly on the worker outcomes (Lazear et al. (2015); Hoffman and Tadelis (2021)). I adopt a different yet complementary approach by identifying the managers that the firm considers of good quality and then looking at their impacts on workers. An advantage of this alternative technique is that it avoids issues of circular reasoning whereby good managers are defined on the same outcomes that are then used to estimate their effects.

[30]The approach in this paper is to study how these managers, who are recognized as particularly productive by the firm, impact their subordinates' outcomes.

shows the correlation with managers' fixed effects in worker pay in logs. Managers' value added measures of this type are the ones most commonly used by the literature (Lazear et al., 2015). I take worker pay five years after the manager exposure to consider the evolution of a worker careers and perform an Empirical Bayes shrinkage procedure to account for the upward bias in the variance due to sampling noise (Morris, 1983).[31] Second, Table 1.4 shows that the high-flyer manager status is positively correlated with a number of ex-post performance measures: managers' future salary growth, probability of promotion to work-level 3, performance appraisals, and workers' anonymous upward feedback on the managers' leadership.[32]

Table 1.5 shows additional descriptive characteristics. High-flyers are more likely to have been developed internally: they are 18ppt less likely to be mid-career recruits and 33ppt more likely to have been hired through the graduate program at the company. In terms of demographics, high-flyer status is positively correlated with being female and havingt a degree in economics and social sciences; which is consistent with positive selection into corporate jobs for women and negative selection for those who have a STEM major. Time-use data from Microsoft on a subset of workers - the whole population of a particular division - reveals that high flyer managers dedicate 0.7 more weekly hours in 1-1 meetings with subordinates (a 21% increase).

It is worth highlighting how promotion speed can easily be applied to other contexts as a way to single out talented leaders: the data requirements are not particularly stringent and are not context-specific. Any organization would have a job ladder for workers and workers' age is easily observable and verifiable.

1.4.2 Manager transitions

I leverage the naturally occurring rotation of work-level 2 managers between teams to conduct an event-study analysis following a manager transition. In an ideal experiment, I would randomize workers with different skills to managers of different qualities and then measure the effects on their career progression in subsequent years. As it would be unfeasible for most real-world companies to randomly shuffle their workers and managers, I use managerial rotations across teams as a natural experiment. These rotations generate variation in the manager types that each worker meets and allow for causal identification of manager effects. I only consider the manager transitions that result from the reassignment across teams as part of managerial lateral

[31] See Appendix Figure 1.57 for the density plot of the manager fixed effects. I define high manager value added if it is above the 75th percentile. Results are similar when using the median as the cutoff.

[32] Baker and Holmstrom (1995), using internal personnel records from a service sector firm, note that a prominent feature of the data is "fast tracks": those who advance quickly early on, continue to advance quickly later on.

rotations; I do not include instances when the manager is promoted to a higher position. I identify such exogenous transition events in the data by observing that the new manager assumes responsibility for all employees in the team. I am not considering transitions that result from employee promotions to another team, or employee transfer requests. I only consider the first manager transition that a worker experiences in the data and then keep tracking the worker outcomes irrespective of the presence of future manager transitions.

These manager transitions are not literally decided by a coin toss, but anecdotal evidence suggests that they are exogenous to workers and teams. Testimonies from executives and HR representatives suggest that these transitions are orthogonal to employee characteristics. As part of corporate strategy, work-level 2 managers are expected to gain experience in different projects and teams within a given sub-function. For this reason, managers are reassigned laterally across teams in random order to gain exposure to different teams and activities and hence broaden their managerial skills. The aim is for the managers to eventually experience all teams within a sub-function. The rotations are also used as a screening mechanism to evaluate who should progress further to work-level 3 (director level). The firm has been implementing this rotation policy for several decades. New assignments tend to occur between 15 and 30 months in the manager's previous position (Appendix Figure 1.58 plots the full CDF of the duration in the previous job). Overall, 74% of managers make at least one of these transitions in my panel data.

Rather than relying exclusively on testimony that these manager rotations are orthogonal to the workers' characteristics, I test this assumption by examining the parallel trajectories of employees who undergo different transitions along a wide range of outcomes using the event-study analysis (see next sub-section 1.4.3 for more details). I conduct additional endogenous mobility tests where I show that an array of past team characteristics in the three years before the manager transition - including team performance, inequality, transfer rates, and team diversity - cannot predict the quality of the incoming manager. To evaluate the correlation between current team characteristics and high-flyer status of future managers, I estimate the following model at the team level:

$$y_{team,t} = \alpha_0 + \pi_0 \; High - flyer \; manager_{team} + \boldsymbol{X'_{team,t}\beta} + \epsilon_{team,t} \tag{1.1}$$

where *High-flyer manager*$_{team}$ denotes the quality of the future manager and controls ($\boldsymbol{X_{team,t}}$) include function, country and year FE, and team size. Under the null of $\pi_0 = 0$ managers cannot impact team performance before they take charge, thus any correlation between change in manager type and past team characteristics is indicative of sorting. Table 1.6 - Table 1.8 show the results: there is no evidence of high-flyer managers being assigned to teams with worsening performance or teams that improve prior to their arrival.[33]

An example can illustrate the empirical strategy. Consider two teams, each supervised by a low-flyer manager. One of these teams then transitions from the low-flyer manager to a high-flyer manager, while the other team transitions from the low-flyer manager to a different low-flyer manager. I compare the outcomes of the workers each month leading up to the manager transition date and each month after the transition. As both teams are affected by a manager transition, this design nets out the effect of the transition on outcomes. I only consider the first manager transition that a worker experiences.[34]

Since the identification strategy relies on manager transitions, I run a similar model as in equation 1.1 but allow for different transitions to have a different impact, leaving the *LowtoLow* transition as the omitted category:

$$y_{team,t} = \alpha_0 + \tau_1 E^{LtoH}_{team} + \tau_2 E^{HtoL}_{team} + \tau_3 E^{HtoH}_{team} + \boldsymbol{X'_{team,t}\beta} + \epsilon_{team,t}$$

In particular, I am interested in testing the hypotheses that $\tau_1 = 0$ and that $\tau_2 - \tau_3 = 0$. Table 1.9 - Table 1.11 show the results and there is no evidence that the type of manager transition is correlated with teams' prior performance.[35]

The event-study data comprises 27,711 transition events, involving 27,711 unique workers and 13,755 unique managers.[36] Events occur every year but are twice as likely to occur in the first two years of the panel (2011-2012) since I only consider the first manager transition. They affect workers in every function and country. The type of event is largely unrelated to the characteristics of the worker, as shown in Table 1.12. The table also shows that whether the worker has an event is negatively correlated with the age and the experience of the worker, as the rotation policy is directed at employees in work-level 1.

1.4.3 Event study design

Let y_{it} be the outcome of interest, where the subscripts i and t denote employees and time, respectively. The main outcomes in my analysis are the employees' number of promotions, lateral transfers, firm exit, and performance such as salary and sales bonus.

The research design is as follows. I compare two teams, each led by a low-flyer manager.

[33] The statistically significant coefficient in Table 1.6 (Column 5, at 10% level significance level) can be due to chance as I am testing 13 hypotheses, and hence there is a 75% chance of observing at least one significant result.

[34] In sub-section 1.5.5, I show that my results are robust to only considering new hires, for whom I can tell for certain that this is their first manager change at the firm.

[35] The statistically significant coefficient in Table 1.9 (Column 5, at 5% level significance level) can be due to chance as I am testing 13 hypotheses, and hence there is a 49% chance of observing at least one significant result.

[36] As I only consider the first transition event experienced by a worker, the number of unique workers is the same as the number of transition events.

One team transitions to a high-flyer manager, and the other team transitions to a different low-flyer manager. Similarly, I compare two teams, each led by a high-flyer manager, where one team transitions to a low-flyer manager, and the other team transitions to a different high-flyer manager.

I specify the model below:

$$y_{it} = \sum_{j \in J} \sum_{s \neq -1} \beta_{j,s} D^{j}_{i,t+s} + \boldsymbol{\xi_t} + \boldsymbol{\alpha_i} + \epsilon_{it} \tag{1.2}$$

where s indexes the months relative to a change in manager and D^j denote the event-study indicators for the periods leading up to and following a transition event $j \in \{LtoH, LtoL, HtoL, HtoH\}$. For instance, $LtoH$ denotes a transition from a low- to a high-flyer manager. ξ_t comprises of year-month FE and α_i is worker FE to control for permanent differences in worker productivity[37]. Standard errors are clustered at the manager level.

The event-study window spans from 36 months before the event to 84 months after the event. The time window is due to the length of the panel data. In particular, because I only look at the first manager transition, most events occur in the first three years of the panel (2011-2013) and hence this constrains the length of the pre-event time window. For example, the -12th quarter estimate is the average of the estimates in months -36, -35, and -34 before the event; only workers that experience the event after December 2013 can identify these coefficients. The omitted category in the leads and lags of the event indicators is the month prior to the event. In the event-study graphs, I aggregate the monthly coefficients to the quarterly level for ease of presentation. I use the interaction-weighted estimator to avoid contamination from effects from other periods (Sun and Abraham, 2020).[38]

Since some outcomes are count variables, such as the number of salary increases and the number of transfers, I also estimate the model in equation 1.2 using a Poisson quasi-maximum likelihood model[39]:

$$E(y_{it}|\boldsymbol{X}_{it}) = exp\Big(\sum_{j \in J} \sum_{s \neq -1} \beta_{j,s} D^{j}_{i,t+s} + \boldsymbol{\xi_t} + \boldsymbol{\alpha_i} + \epsilon_{it} \Big) \tag{1.3}$$

[37] The worker fixed effects also account for different starting points (initial age or workforce experience) and the time fixed effects then account for the variable increasing by the same amount for each worker.

[38] In practice, heterogeneous results are unlikely to bias the results. First, the TWFE estimates are very similar to what is obtained when using the (Sun and Abraham, 2020)'s estimator. Second, I implement a test for the potential influence of negative weights proposed by (De Chaisemartin and d'Haultfoeuille, 2020). I find that the total sum of the negative weights is always < -0.01 for all outcomes. The contamination from other treatments is similarly small (< 0.06). Because all weights must sum to one, these results indicate that the negative weights and contamination from other treatments are not influential in this setting.

[39] The estimator is consistent in the presence of high dimensional fixed effects and can be used to model non-negative dependent variables without the need to specify a distribution (Correia et al., 2020).

To isolate the impact of a change in manager type from a change in manager more generally, I always compare employees undergoing manager transitions where one of those transitions results in a change of manager type and the other does not. Hence, the estimates of interest are the differences between types of transitions: $\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s}$ (i.e., transitioning from a low-flyer manager to a high-flyer manager, relative to transitioning from a low-flyer manager to another low-flyer manager) and $\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s}$, where s indicates the time since (or until) the transition date.

The key assumption is that, prior to the transitions, employees were on the same career trajectories irrespective of their upcoming transition. The event-study framework provides a natural test of the identifying assumption: I can assess the evolution of the outcome in each month before the date of the transition to confirm if the trends were truly parallel before the event date.

1.5 Results

In this section, I document the effects of gaining a high-flyer manager on the workers' lateral and vertical moves, exit from the firm, and career progression. I discuss the transition in the opposite direction, i.e. losing a high-flyer manager, in Section 1.6; where I also provide additional supporting evidence on the job-allocation margin being the key factor underlying the observed impacts of high-flyer managers.

1.5.1 Workers' transfers and exit from the firm

Figure 1.4 presents the effect of gaining a high flyer manager based on the econometric model discussed in Section 1.4: it compares the effects on the number of lateral moves when transitioning from a low to a high-flyer manager (*LtoH*) relative to transitioning from a low manager to another low-flyer manager (*LtoL*). This event-study graph shows the evolution of the number of lateral moves in each of the 12 quarters (3 years) leading up to a manager transition and the 28 quarters (7 years) after the manager transition. The quarter before the event (-1) corresponds to the omitted category, and thus the corresponding coefficient is always zero by construction.

Figure 1.4 shows that, prior to the event date, the differences in the coefficients are statistically indistinguishable from zero. This evidence indicates that the assumption about parallel trends holds. In contrast, after the transition date, the evolution of lateral moves starts to gradually diverge between the *LtoH* and *LtoL* workers. The moves increase up to 20 quarters after the manager transition and then level off at the new higher level. At 28 quarters after the manager transition, the lateral moves are 0.15 higher (or a 30% increase, p-value <0.05).[40]

In addition, I isolate task-distant lateral transfers. First, I consider cross-functional moves, that involve a major horizontal job change, such as from HR to R&D.[41] Figure 1.5 shows that these also increase by 0.07. Second, in Figure 1.6, I match the MNE job titles to the Occupational Information Network (O*NET) job classification data, which provides different scales on skills and activities required for each job, and construct angular separation measures of task distance across jobs (Gathmann and Schönberg, 2010). The O*NET data produces multiple scales of job descriptors such as work context, work activities, abilities, and skills. My baseline specification uses the skills vector but the results are robust to taking the average of the different distance measures.[42] Both Figure 1.5, which looks at cross-functional transfers, and Figure 1.6, which looks at task distance in transfers based on O*NET data, paint a consistent picture of an increase in task-distant transfers (which represent an increase of around 20%).

I can decompose the overall increase in lateral transfers by whether they occur within the team, outside of the team but within the same function, or across functions. Figure 1.14 takes the event study coefficient at the 8th quarter (approximately at the end of a manager rotation, as they last on average two years)[43] shows that half of the job moves are within the team, 38% are outside the team but within the same function and the remaining 12% are across functions.

Appendix Tables 1.13 and 1.14 show the transition matrices across functions, for the "low-to-high" transitions and the "low-to-low" transitions respectively. Moves can occur across most of the possible function-to-function combinations but the biggest moves are from Customer Development (sales) to Marketing, from Human Resources to General Management or Marketing, and from R&D to Supply Chain. In addition, I zoom-in into two functions, Supply Chain (Tables 1.17 and 1.18) and R&D (Tables 1.19 and 1.20), to look at moves across sub-functions but within a function. Overall, in both cases, it can be noted that there is a lower share of workers remaining in the same sub-function (lower diagonal values) and a higher share of workers transitioning to other functions (higher values in the penultimate column).

In Figure 1.7, I focus on work-level promotions. These are major promotions that reflect meaningful changes in job responsibility such as transitioning from a work-level 1 front-line worker position to a work-level 2 managerial position. At 28 quarters after transitioning to

[40] Appendix Figure 1.70 plots the probability of making at least one job move since the manager transition and shows that the effects of gaining a high-flyer on number of lateral moves come from many workers making at least one lateral move (an increase of 10ppt or 53% for the workers in the *LtoH* transition with respect to the workers experiencing the *LtoL* transition), rather than few workers making many lateral moves.

[41] There are 14 functions in the MNE; the biggest six are: Sales, HR, R&D, Supply Chain, Finance, Marketing.

[42] I provide more details on the construction of the task distance measure in the Appendix 1.15.

[43] I need to define a reasonably short time window to consider the within team job moves so to evaluate them while the original transitioning manager is still in charge. It is however important to note that the cross-functional transfers take longer to occur and they keep increasing until the 22nd quarter.

a high-flyer manager (relative to transitioning to another low-flyer manager), the work-level promotion rates are 0.06 higher (an increase of 40%, p- value<0.05). These major promotions start to occur at a relatively later stage compared to the lateral transfers, around 6 quarters after the manager transition. Moreover, Appendix Figure 1.62 shows that, conditionally on being promoted to work-level 2, the workers promoted under a high-flyer manager perform better in terms of pay growth and of the anonymous manager rating score given annually by workers.

I also assess whether there is an effect on worker exit from the firm. Figure 1.8 shows that there is no impact on worker exit and Figure 1.9 shows that the results are the same whether I look at voluntary (quits) or involuntary (layoffs) exits.[44] Moreover, Figure 1.29 shows that there are no heterogeneous effects by whether the worker is an under or over-performer in terms of pay growth at baseline. As I find a higher rate of job transfers but no evidence of higher firm exit, this suggests that high-flyer managers are not kicking out lower-performing workers from the firm but rather they are finding alternative suitable deployments inside the organization.

1.5.2 Workers' career progression

I have shown evidence that the high-flyers cause higher job reallocation to the workers they supervise through lateral and vertical transfers. In this sub-section, I show that high-flyer managers also have a positive persistent impact on the workers' career progression. Figure 1.10 compares the effects on the number of salary grade increases when transitioning from a low to a high-flyer manager (*LtoH*) relative to transitioning from a low manager to another low-flyer manager (*LtoL*). Prior to the event date, the differences in the coefficients are statistically indistinguishable from zero. In contrast, after the transition date, the evolution of salary increase rates starts to gradually diverge between the *LtoH* and *LtoL* workers. It keeps diverging up to the 20th quarter after which it levels off at the new higher level. At 28 quarters after transitioning to a high-flyer manager (relative to transitioning to another low-flyer manager), the salary grade promotion rates are 0.25 higher (p- value<0.05).

This corresponds to a salary that is 30% higher: Figure 1.11 shows the salary (pay plus bonus) estimates and the bonus estimates separately, which show an even bigger increase of 125%. Despite the large estimate, it should be considered that bonus is around 10% of fixed pay for work-level 1 workers, and the effect on total pay is mainly driven by the increase in fixed pay as shown in Appendix Figure 1.64.[45] The gap in overall pay is economically large: in the

[44] While I find no differential impact on exit, it is worth noting that the reasons a worker may quit the organization may be different under a high-flyer manager. Appendix Figure 1.63 shows some exploratory evidence from the voluntary survey of workers who quit. The workers who quit under a high-flyer manager are less likely to report cultural fit and line manager as the exit reasons while they are more likely to report a change of career.

[45] The compensation data is only available from 2016 onwards, hence I can estimate the post-transitions coefficients only.

U.S. it represents \$29,373 in annual salary, on average.[46] An alternative way of illustrating the magnitude of this effect is to consider that a 30% higher salary corresponds on average to seven additional years of tenure for a new hire at entry level.

When inspecting Figure 1.10, it is important to keep in mind that these coefficients refer to differences across transition types. As a result, a coefficient of zero in the post-treatment period does not imply that workers remain in the same pay grade; rather, it indicates similar salary growth rates across workers transitioning from low to high-flyer managers versus workers transitioning from low to another low-flyer manager. Workers' salary grades tend to increase over time as this context is characterized by plentiful upward mobility.

Figure 1.13 shows a similar analysis but at the team level, using a 24 months horizon (the average duration of the manager assignments). It confirms that overall team performance increases, as measured by average pay growth, and that there is a higher churning of workers across teams. Performance dispersion in terms of performance appraisals[47] also increases but this comes from the top of the distribution (workers scoring strictly above 100) as there is no change in the bottom share (workers scoring strictly below 81).

1.5.3 Workers and factories' productivity

Do the effects of the high-flyer managers result in higher objective worker productivity or are they leading to the worker earning higher pay for the same performance? So far, I have interpreted higher worker pay growth as evidence of higher productivity. By leveraging objective performance data from the subset of Indian sales workers (circa 2500 workers), I can provide further evidence in favor of this interpretation.[48] The productivity data is obtained from the sale incentives records and it represents the monthly sales bonus in Indian rupees, whereby field sales workers in India are paid a monthly variable sales bonus according to how much they achieved relative to their targets.[49] The data is high-frequency as sales performance is tracked monthly but it is only available for 2018-2021, and it is relatively noisy. I can hence run a similar model to Equation 1.2 using a static version of the event study given the limited time window

[46] To quantify how influential high-flyers are for workers' careers, one can also compute how they affect the present value of the workers' lifetime income. Assuming that careers last another 30 years (since most workers are in their late 20s or early 30s) and using a discount rate of 5% (I follow Frederiksen et al. (2020) for this assumption), a two-year exposure to a high-flyer manager is associated with an increase in the PDV of pay of 375% of average annual pay.

[47] The scoring range is 0-150.

[48] While most of the data come from the global personnel records, sales data is managed independently in each of the countries and the data needs to be separately collected on a country-by-country basis by liaising with the countries' local sales teams. A second data challenge is that the field sales teams are increasingly being outsourced to contractors. India is the country where outsourcing had still not taken place at the time of data collection and it is also the country with the largest number of workers in the MNE.

[49] Some examples of sales targets include: growth of sales; product placement; on-shelf availability; additional exhibitions; the number of orders vs total visits each month.

available:

$$y_{it} = \sum_{j \in J} \beta_j Post^j_{it} + \boldsymbol{\xi_t} + \boldsymbol{\alpha_i} + \epsilon_{it} \tag{1.4}$$

where $Post^j$ denote the indicators for the onset of a transition event $j \in \{LtoH, LtoL\}$, ξ_t comprises of year-month FE, and α_i is worker FE to control for permanent differences in worker productivity. Standard errors are clustered at the manager level. I cannot look at the reverse transition, losing a high-flyer manager, as the observations for the $HtoH$ manager transition are too few in this sales sub-sample.

Figure 1.15 shows that sales performance increases by 27% upon switching from a low to a high-flyer manager. This amounts to an increase of INR 2,650 in monthly sales bonus pay.[50] The next rows show that I can replicate the main findings also in this sub-sample: overall salary increases by 7.9% and the number of lateral transfers increases by 12.7%. Lateral moves in this context generally consist in changing selling products, clients, or geography.

I conduct some additional descriptive exercises that combine the lateral moves with the performance effects. I separate the sample of workers who make at least one lateral move after the manager transition - up to five years after the manager transition - from the workers who do not. For the workers who make at least one move, I compare the within-worker change in sales performance between those who are exposed to a high-flyer (as a result of the managers' rotations) and those who are exposed to a low-flyer. For the workers that do not move, I compare the sales performance *ex-ante*, before the manager transition, between those who are exposed to a high-flyer and those who are exposed to a low-flyer, again because of the managers' rotations. Figure 1.16 shows the estimated coefficients: for workers that move, those who do so following a high-flyer manager experience a 71% improvement in sales performance (post-manager-transition) compared to the workers who move after being exposed to a low-flyer. For the workers that do not move, those that experience a high-flyer manager transition after were 34% more productive *before* the manager transition. This suggests that transfers *per se* do not necessarily have a positive impact on productivity and high-flyer managers are generating the right transfer for the right worker, or better *worker-job matches.*

Altogether, an efficient allocation of workers to jobs within sites should reflect into higher productivity in the entire site, creating a link between individual-level effects and firm-level outcomes. I provide suggestive evidence, at a correlational level, that factories with workers that had more past exposure to high-flyer managers are indeed more productive on the whole. I obtain a measure of productivity at the factory-year level, *output per worker*, and a measure

[50] According to the currency exchange rate on October 25th, 2022, this would be \$32 where \$1 = INR 83.

of costs per unit of output, *costs per ton*, for all factories globally over 2019-2021.[51] For each worker, I construct a measure of past exposure to high-flyers as the share of months with a high-flyer up to the year before productivity is measured (one-year lag). This evidence is only correlational in nature as the variation at the factory level in the workers' past exposure to high-flyers is not necessarily exogenous.

I regress output per worker in logs against workers' past exposure of high-flyer managers in the factory, clustering the standard errors by factory-year.[52] The regression controls for country, product category and year fixed effects, the share of managers, and the number of blue-collar and white-collar workers at the factory.[53] Figure 1.17 shows that increasing workers' past exposure to high-flyers by 10ppt is associated with an increase in output per worker by 20%, that is the semi-elasticity between the two variables is equal to 2.03. Similarly, the semi-elasticity between costs per ton and workers' past exposure to high-flyers is -1.4. Altogether, these two results indicate that the high-flyers' effects are increasing profits, taking prices as given.[54]

To further probe the mechanism of high-flyers leading to more productive worker-job matches, I look at the relationship between the workers' number of job moves and factory productivity, and how that depends on workers' previous exposure to high-flyer managers. I separate lateral moves into two groups depending on whether they occur after the worker is exposed to a high-flyer or a low-flyer manager (up to five years after the manager transition). In Figure 1.19, I regress output per worker against the number of job moves, with both variables measured in logs. While I find a significant positive impact on productivity for lateral moves originated by high-flyer managers, the slope is flat for those originated by low-flyer managers. Figure 1.20 shows a similar pattern for costs per ton; they decrease (slightly increase) the higher the number of lateral moves induced by high (low)-flyer managers. These results echo the finding from workers' sales bonuses: transfers *per se* do not necessarily have a positive impact on productivity. It is the lateral moves induced by high-flyer managers that bring about higher productivity while the moves that follow from low-flyer managers have zero or even negative impact on productivity (albeit not statistically significant, the slope for costs per ton is positive).[55]

[51] Tons of products produced per FTE is a common KPI for manufacturing firms (FTE stands for Full-Time Equivalent). The cost per ton measure considers the operational costs per ton which are predominantly made up by labor and energy costs (it does not include the cost of raw materials). Because of changing reporting requirements at the firm, the costs per ton data could only be shared for the main product category (there are three product categories in total).

[52] Appendix Figure 1.69 shows the same regression using the share of high-flyer managers as the independent variable and Appendix Figure 1.68 shows that there is variation in the share of high-flyer managers across factories.

[53] Factories tend to have a non-trivial share of white-collar workers that manage operations and blue collars. The overall share of white-collar workers in factories is 24%.

[54] The assumption on constant prices is plausible for two reasons: first, I am controlling for country, product category and time fixed effects in the regression, and second, pricing is the responsibility of marketing teams and production managers do not set prices.

[55] Moving workers around could in fact be detrimental for productivity if there is no meaningful improvement

1.5.4 Workers' behavioral and attitudinal data

To further shed light on the channels behind the results, I make use of three additional sources of data, the career mobility platform, a platform for flexible short-run projects, and the engagement annual surveys. They complement the administrative data in unveiling more concretely how worker behavior may change when supervised by a high-flyer manager. For these sources, I do not have a window long enough to run an event study.[56] Hence, I estimate the following static model on the switchers' sample (after the manager transition) to estimate the contemporaneous impact of having a high-flyer manager:

$$y_{it} = \alpha_0 + \alpha_1 High\ flyer\ Manager_{it} + \boldsymbol{X_{it}}'\boldsymbol{\beta} + \eta_{it} \tag{1.5}$$

where the coefficient of interest is α_1 and $\boldsymbol{X_{it}}$ controls for the manager's age group, tenure, and tenure squared interacted with the manager's gender.

Figure 1.21 shows that workers gaining a high-flyer manager are more likely to engage in the career mobility platform and become "active learners", defined as having posted at least three focus skills, completed at least five courses/items, and shared at least one item with a colleague.[57] This is an increase of 13% relative to a low-flyer manager. As this platform is used as an internal labor market tool to combine skill acquisition and career mobility, these results complement the findings on the higher lateral transfer rates.[58]

Moreover, Figure 1.22 looks at worker participation in the initiative of the flexible projects, which was conceived to allow greater career and organizational agility by empowering workers to design their own career paths. Workers gaining a high-flyer manager are 17% more likely to register on the platform of the flexible projects (6ppt), 6% more likely to complete their profile in full (3ppt), 12% more likely to state that they are available for flexible opportunities (3ppt), 87% more likely to report being available to be a mentor (11pt), and 24% more likely to apply for jobs (1ppt). Although launched in mid-2018, the flexible project program has taken some time to gain momentum, also due to the COVID-19 pandemic. Hence, the baseline take-up is still low with only 2.3% of workers applying for flexible projects, but it is projected to improve in the future.

Figure 1.23 shows that workers under a high-flyer manager are also more likely to report

in job match as it could also lead to previously accumulated job-specific human capital remain unused. The framework in Section 1.7 clarifies this trade-off.

[56] The career mobility and flexible projects platforms were established in 2018 and the first global annual pulse survey was run in 2018.

[57] The reporting thresholds are determined by the platform.

[58] Appendix Figure 1.50 shows the skills sought after by the employees at the firm. It is interesting to note that the biggest skill in demand by far is management, underscoring that good managers are scant.

higher manager (6%) and team effectiveness (2.4%), while I do not find statistically significant differences in job satisfaction (albeit the coefficient is positive and far from zero), autonomy and company effectiveness. This backs the intuition that, if workers are better matched to their day-to-day to tasks, team performance and efficacy increase. As the survey contains many variables, I aggregate them together in four indices, by grouping together the variables by theme, and taking the first principal component (results are unchanged if I take a simple average, see Appendix Figure 1.65).[59] Moreover, since most people tend to answer three or four out of the 5-point Likert scale, I use binary indicators for whether a worker reports five.

1.5.5 Robustness

Restricting the event-study to a single cohort

Since my panel covers 132 months, there is a mechanical restriction on the workers that identify the medium-run effects in the event study. That is, since the 28-quarter estimate is the average of the estimates in months 82, 83, and 84, only workers that experience the event before January 2015 can identify these coefficients. Even for workers who are in the panel in all periods, these coefficients are identified only from events that occur before January 2015. I show that these composition effects do not drive my results by replicating the analysis on a single cohort of workers. I restrict the workers who experience an event to those who have it before January 2015. I retain 68% of the workers who experience a transition event of any kind. Appendix Figures 1.74 - 1.78 show that the event studies limited to this cohort of workers retain the timing and magnitude of the baseline results.

Restricting the event-study to new hires

Throughout the paper, I am only considering the first observed manager transition. However, as my data is only available from January 2011, some workers may have experienced other manager transitions before then. If so, my estimates are averaging the effects on workers that have different histories in terms of manager transitions. This should not cause bias in my estimates as long as each transition event is independent, which follows from the natural experiment. However, I am underestimating the effect of the first manager transition from low to high-flyer as some of these workers may have had additional high-flyer managers in the past. I show that my results are robust to only considering new hires, for whom I can tell for certain that this is their first manager change at the firm. I retain 52% of workers, who I observe in the data since they enter the firm. Appendix Figures 1.79- 1.83 show that the event studies limited to new

[59] Refer to Appendix Table 1.24 for the list of variables and their grouping in themes.

hires retain the timing of the baseline results and, as expected, the estimates are larger.

Poisson model for count data

Appendix Figures 1.71 - 1.73 show the event-study graphs when using a Poisson model as in equation 1.3 for the count variables: salary grade increases, lateral and vertical transfers. The figures report the first differences in the exponentiated coefficients and so they should be interpreted as the difference in the incidence rate ratios. For example, Figure 1.73 indicates that, 5 years post-transition, workers gaining a high-flyer manager have a rate of salary increases 1.3 times greater.

Placebo events: managers' position number oddness

As a robustness check, I reproduce the analysis, but instead of focusing on high-flyer managers as the relevant characteristic of managers, I focus on a characteristic that I know ex-ante should not be relevant: whether the manager's "position number" (generated automatically by the HR system when hiring a worker) is even or odd.[60] This placebo test provides a useful sanity check. First, it helps rule out mechanical reasons why my event-study framework would generate spurious effects. Second, this placebo analysis can be used to assess whether my standard errors are adequate: e.g., if I found statistically significant coefficients, it would suggest that the inference is misleading.

The regression of interest is identical to the main specification from equation 1.2, except that managers' success is replaced everywhere by the managers' position number oddness. Hence the set of manager transitions can be denoted as $j \in \{EtoO, EtoE, OtoE, OtoO\}$. I identify analogous difference estimates for these placebo events. For example, the difference estimate $\hat{\beta}_{EtoO,s} - \hat{\beta}_{EtoE,s}$ measures how workers react to gaining an odd-number manager (i.e., transitioning from an even-number manager to an odd-number manager, relative to transitioning from an even-number manager to another even-number manager).

Appendix Figures 1.84-1.87 are equivalent to Figures 1.4-1.10, but they are based on manager's position number oddness instead of high-flyer status. As expected, they show no significant difference between the two types of transition, either before or after the event. For instance, at 10 quarters after transitioning from an even-number to an odd-number manager (relative to another even-number manager), the difference between the number of pay grade increases of odd-number and even-number workers is very close to zero, statistically insignificant, and

[60] The position number is distinct from the employee ID number, the official number used for identification of an employee inside the firm. The position number is also unique at the employee level but it is only used administratively by HR.

precisely estimated.

1.6 Discussion

1.6.1 Evidence for the matching channel

The results in Section 1.5 show higher transfers and career progression for workers gaining a high-flyer manager. I provide evidence indicating that matching workers to jobs is a key mechanism underlying the observed impacts of high-flyers on workers' careers.

Mediation analysis. To further analyze the role of lateral moves behind the increase in salary, I perform a mediation analysis following the method by Imai et al. (2010a) and Imai et al. (2010b). The underlying intuition is that the treatment effect of high-flyers on outcome Y can be decomposed as operating through the mediator M:

$$\frac{dY}{d\text{High-Flyer}} = \frac{\partial Y}{\partial M}\frac{\partial M}{\partial\text{High-Flyer}} + R \tag{1.6}$$

where R is the part of the treatment effect which cannot be attributed to the mediator. The actual implementation is based on an algorithm that calculates the average mediation and direct effects by simulating predicted values of the mediator or outcome variable, which are not observed, and then calculating the appropriate quantities of interest (average mediation, direct effects, and total effects). The results of the mediation analysis should be interpreted with caution and to be able to interpret it causally one needs to make strong assumptions about the source of variation of the mediator. As I lack exogenous variation in the transfers and have to rely on a single source of exogenous variation (the rotation policy), this analysis can only provide suggestive evidence of the importance of the mediator in explaining the treatment effect. Despite these limitations, the exercise can complement the evidence of whether the job-matching mechanism explains the high-flyer managers' effects.

I take the number of salary grade increases in the 28th quarter as the outcome, Y, and the number of lateral moves in the 20th quarter as the mediator, M.[61] I find that lateral transfers contribute to 62% of the total effect of high-flyers on the number of salary increases. The results are very similar when using the approach by Gelbach (2016) and Heckman and Pinto (2015).

Taken together, this evidence paints a consistent picture of the role of transfers in mediating the impact of high-flyer managers, but there are good reasons to believe that 62% is a lower bound of the importance of the job matching channel. By using transfers as the instrument

[61] Results do not change for small changes to the time horizons.

to proxy for job matching, the analysis misses the gains of (i) workers that do not change jobs because they are in good matches already, (ii) vertical transfers (which are also about job allocation but are left out as they involve a salary raise by definition), (iii) any task-allocation decision that does not involve a job change, such as the assignment of short-term projects.

Asymmetric effects for losing a high-flyer manager. So far, I discussed the evidence of what happens upon gaining a high-flyer manager. I now look at the reverse transitions, i.e. losing a high-flyer manager (moving from a high-flyer to a low-flyer manager compared to moving to another high-flyer manager). Figures 1.30-1.37 show that there is no differential impact in losing a high-flyer manager, the estimates are close to zero and statistically insignificant. Since only 9% of the events are $HightoHigh$, by virtue of the definition of a high-flyer manager[62], it should be kept in mind that these results are less conclusive than those for gaining a high-flyer manager. Due to the smaller sample size, the confidence intervals are wider, and especially the coefficients leading up to the transition are more imprecisely estimated.[63] Yet, the point estimates are clearly smaller compared to Figures 1.4-1.10 and also do not exhibit a detectable downward trend, as it would be expected if losing a high-flyer manager had a symmetric effect of gaining a high-flyer manager. Hence, the high-flyer manager results are asymmetric: compared to gaining a high-flyer, losing a high-flyer does not lead to similar findings in the opposite direction such as lower salary growth and transfers (see Figures 1.39 - 1.43 for a formal test of asymmetries).

This evidence illustrates that there are dynamic benefits of a one-time exposure to a high-flyer manager[64]: the impact endures even after transitioning to a low-flyer and there is no additional impact of having a second high-flyer manager. These findings reinforce the interpretation of the allocation channel as, once a worker has found the right job match, the gains cannot be erased by transitioning to a low-flyer manager. If high-flyers were mainly motivating or monitoring workers to exert higher effort, we would expect to see symmetric effects so that, upon transferring from a high- to a low-flyer manager, there is a negative impact on the worker's career progression (compared to transferring from a high- manager to another high-flyer manager).

There is one caveat to bear in mind when comparing the impact of gaining versus losing a high-flyer manager. Unlike the manager transition used in the identification strategy, the *first* manager-worker assignment is not necessarily random. In fact, the identification strategy relies on the *second* manager-worker assignment being orthogonal to worker characteristics, but not

[62] As a reminder, the share of high-flyer managers is 30%.

[63] As noted before, most of the transition events occur in the first three years of the panel (2011-2013) since I only consider the first manager transition. For example, the -12th quarter estimate is the average of the estimates in months -36,-35, and -34 before the event; only workers that experience the event after December 2013 can identify these coefficients.

[64] The average duration of a manager assignment to a team is two years.

necessarily the first assignment, which may be a result of sorting. In practice, it is impossible to check for this given the data is only available for 2011-2021 (i.e. I cannot observe the workers' histories before 2011). Hence, any differences in the outcomes of workers that start with a low-flyer manager against those of workers that start with a high-flyer manager could be either due to the treatment effect of high-flyers while managing the workers on the job or due to differential selection of workers by manager type ex-ante. One could for instance imagine that the ability of high-flyers to spot unique talents occurs even before interacting directly with the worker in the day-to-day job, at the interview/selection stage. Overall, one might view this caveat as less critical for the validity of my results (relative to other settings) given that managers being able to select the right workers for their team is highlighted as the key channel that differentiates high-flyer managers from the rest.

Task-distant transfers. Alongside the improvements in the workers' career progression, I document an increase in the rate of lateral job transfers, including task-distant ones. I define task-distant moves in two ways; based on transfers across functions and based on O*NET task distance measures across jobs (Figure 1.5 and Figure 1.6). This evidence on higher task-distant moves cannot be easily reconciled with other channels such as high-flyer managers engaging in better teaching or transmitting higher motivation to work.

Pay inequality. I complement the worker-level regressions with team-level analysis to look at pay inequality. To shut down effects due to changes in team composition, I keep the team constant at the time of the manager transition, regardless of whether a worker continues to be working under the manager of the transition or changes manager after some time. For each team and month, I compute the coefficient of variation in pay among the team members between 36 and 60 months after the manager transition, averaging the monthly estimates to quarters. Again, the estimates of interest are the differences to isolate the impact of a change in manager quality from a change in manager more generally.

I find that teams transitioning from a low- to a high-flyer manager experience a higher coefficient variation in pay (an 18% increase at 28 quarters relative to teams transitioning to another low-flyer manager, Figure 1.12). The increase in dispersion suggests that these managers are identifying the most talented workers and exacerbating the natural differences in ability by directing them to the jobs most suited to the workers' skills. This is another result that would be inconsistent with managers only engaging in teaching the workers, which would predict a *lower* variance in performance among team members. However, the differentiation among talents is not solely vertical. I show next that even the lowest performer at baseline experiences a higher salary growth after transitioning to a high-flyer manager (and higher transfer rates).

New hires. I test whether new hires perform better if hired by a high-flyer manager. While managers have little personal interaction with applicants at the hiring stage, they typically screen CVs and conduct a half an hour interview. Under the hypothesis that high-flyers are better at identifying good candidates for a specific job, their hires should be more positively selected compared to those of low-flyers. Figure 1.26 looks at pay four years after the worker is hired, and provides support to the hypothesis and shows that new hires perform better when hired by a high-flyer manager and that the performance gap is the same irrespective of the next manager transition due to the rotation policy.

Moreover, in Appendix Figure 1.89, I show similar effects when focusing on workers in the graduate scheme designed for fresh college graduates. During the scheme, workers rotate teams every year for three years and at the end of the program they can either be fired, or be allocated to a job inside the firm under a standard contract. I show that exposure to at least one high-flyer manager during the program has a positive impact on worker performance. I look at workers' salary in the two years after the program. The effects are stronger for workers that move to a job supervised by a low-flyer manager compared to workers that move to a job supervised by a high-flyer manager, repeating a similar pattern found in the asymmetries in the results (i.e. it is enough to be exposed to a high-flyer once).

Heterogeneous effects. I extend the model in equation 1.2 to test for heterogeneous treatment effects, allowing for heterogeneity in H_i:

$$y_{it} = \sum_{j \in J} \sum_{s \neq -1} \beta_{j,s} D^j_{i,t+s} + \sum_{j \in J} \sum_{s \neq -1} \beta^H_{j,s} D^j_{i,t+s} \times H_i + \boldsymbol{\xi_t} + \boldsymbol{\alpha_i} + \epsilon_{it} \tag{1.7}$$

where all the variables are defined as in equation 1.2. Let H_i be a dummy variable that indexes for example larger offices, then β identifies the effect of high-flyers on small offices while β^H identifies the differential impact between large and small offices. Thus, β^H tests for the presence of heterogeneous treatment effects and it is the main coefficient of interest. Since the high-flyer managers appear to have the largest impact on worker outcomes in the 20th quarter, the display of the heterogeneity analysis focuses on worker heterogeneous outcomes (β^H) in that quarter.

I explore a number of dimensions of heterogeneity. First, I look at workers and managers' characteristics: manager tenure, manager sharing the gender with the worker, the manager and the worker being in the same office, worker age, and worker tenure. Second, I consider characteristics concerning the environment in which they operate: office size, number of different jobs in the office, gender norms in the country[65] and country labor laws.[66] Third, I look at worker baseline performance in terms of average pay growth in the two years preceding the manager transition: above the median, top 10 versus bottom 10 and average team performance. Finally, I consider two endogenous variables to clarify the mechanism: whether the worker changes job and/or manager within two years since the manager transition.

Figure 1.27 shows that the effects are strongest for (a) managers with higher tenure, (b) workers that are in the same office as their manager, (c) younger workers, (d) workers with lower tenure while (e) there is no differential impact on workers that share the same gender with the manager. Conditional on having a high-flyer manager, a higher manager tenure in the firm tends to correlate with more information regarding job opportunities and career paths at the firm, as well as with higher general experience in managing workers. Second, being in the same office as one's manager facilitates interactions and observation by the manager. The larger effects for younger and less experienced workers make sense when thinking that these workers have just started operating in the labor market and they have much to discover about their skills and fit. The gender result indicates that there are no heterogeneous effects along this dimension. I expand on this result in subsection 1.6.2.

Figure 1.28 shows that the gains are larger for (a) bigger offices, (b) offices with a larger number of different jobs, and (c) countries with stricter labor laws. The heterogeneous effects along these dimensions provide further support of the matching channel: small offices or offices with a fewer number of different jobs have less job variety and hence there is less scope for reallocation, and stricter labor laws impose constraints on hiring and firing making reallocation of existing talent to jobs particularly crucial.[67] The last panel plots the gender gap in outcomes (women minus men) and how it varies by the female over male labor force participation in the country as an indicator of gender norms. It shows that high-flyers are more relevant for the careers of women in countries with stricter gender norms. In Ashraf et al. (2022), we analyze the same data and we show that women are positively selected in countries with stronger gender norms and that firm productivity would increase by 32% if we could equate the labor force participation across genders. Hence, it pans out that high-flyer managers have a stronger positive

[65] I use the ratio of the female over male labor force participation in the country. The data is taken from the World Bank.

[66] I use the Restrictive labor Regulations Index from the World Bank. It is available for the period 2008-2020 and it is based on an annual survey of the most problematic factors for doing business (e.g. corruption, taxes, inflation, etc.). The survey is administered to a representative sample of around 15,000 business executives in 150 countries. The Restrictive Labor Regulations Index includes measures related to labor-employer relations, wage flexibility, hiring and firing practices, performance pay, labor taxes, attraction, and retention of talent.

[67] The heterogeneous effects by labor laws echo the findings of Fenizia (2022) on managers having large impacts on the efficiency of the public sector despite the lack of many of the tools available to private sector firms such as hiring, firing, and promotions.

impact on women's careers as opposed to men's in countries with stronger gender norms.

Figure 1.29 displays no evidence of heterogeneity in different dimensions of worker and team performance. I construct the average pay growth for each worker in the two years before the transition and I define whether a worker was above or below the median. I also compare the top 10 percent of workers against the bottom 10 percent of workers. In both of these cases, I do not find clear evidence of heterogeneous effects. This indicates that high-flyers are not disproportionately benefiting higher or lower-performing workers. For instance, the 10th percentile split shows that the weakest workers' career progression (i.e. the bottom 10 percent) also improves when transitioning from a low- to a high-flyer manager, compared to a worker of similar past performance but transitioning to another low-flyer manager. The outcome that shows some more signs of heterogeneous effects is the lateral moves, where the point estimate tends to be lower for higher-performing workers at baseline. This makes intuitive sense when considering that there is a lower chance to find a better job match if a worker is already been performing well in the current job. Panel (c) of Figure 1.29 also shows no heterogeneous effects for baseline team performance.[68]

Panels (d) and (e) of Figure 1.29 show that workers who change manager or job within 2 years of the transition, do not have differential effects. Specifically, the career gains upon moving to a high-flyer manager are not exclusively coming from workers who remain with the same high-flyer manager (panel d) or from the workers who undergo a job change (panel e). The estimates in panel d are useful to cast aside a purely "plug-in" channel of high-flyers, whereby the workers experience a higher career progression by remaining around these managers throughout their career.[69] The fact that there is no differential impact for workers who change and do not change jobs (panel e), also shows that the performance gains are not due to an effect of transfers by themselves, for e.g. a direct training effect of transfers.[70] They are due to high-flyer managers allocating talent to jobs efficiently (i.e. it is about the *worker-job match*, not the *transfer per se*).

[68]It is helpful to consider this result in light of the identification strategy that relies on manager rotations. A threat to the validity of the strategy is the non-random assignment of managers to teams. A profit-maximizing firm may want to design rotations so to maximize output, which may cast doubt on the firm's rationale for having rotations in random order. In light of my results, the best policy would be to match managers and teams so to maximize the chance that each worker gets exposed to a high-flyer manager at least once. This is because the results show that a one-time exposure to a high-flyer has a persistent effect on a worker's career. While this policy is different from moving managers to teams in a quasi-random fashion, it is also different from a policy that dictates positive or negative assortative matching. Overall, there are also a lot constraints the firm operates under such as average length of assignment of two years, availability of vacancies and the requirement of exposing managers to every team within a sub-function.

[69]Appendix Figure 1.88 also shows that workers gaining a high-flyer manager are not more likely to have another high-flyer manager as the next manager.

[70]One might think that doing more job transfers leads to better-trained workers that can then perform better and are more likely to be promoted in the future, irrespective of a worker-job matching channel.

1.6.2 Ruling out other channels

In this section, I consider a number of alternative mechanisms and show that, although these may well exist in the firm, they do not explain my results.

Manager bias. I interpret the results on the workers' career progression as reflecting the causal impact of high-flyer managers on worker performance. The primary threat to this interpretation is that, rather than increasing worker performance, managers are increasing workers' pay without any increase in worker productivity. An extreme view could argue that the results found are due to high-flyer managers inflating their workers' pay and promotion prospects, because of having leniency bias with respect to their workers for instance (Frederiksen et al., 2020). It is important to note that, for this interpretation to hold, the leniency bias must be correlated with the high-flyer manager status. Otherwise, in the case that leniency bias is present but is uncorrelated with being a high-flyer manager, it would be shut down by design as my methodology compares worker outcomes across different types of transitions. Moreover, I present three pieces of evidence indicating that manager bias is unlikely to drive the estimated effects of high-flyer managers on workers' careers.

First, having a high-flyer manager causes higher worker objective productivity as shown in Figure 1.15: being exposed to a high-flyer manager increases monthly productivity by 27%.

Second, I do not find that the workers' lateral and vertical moves occur within the managers' networks, ruling out explanations related to high-flyer managers having greater social connections within the MNE. Figure 1.24 shows the coefficients obtained by running a static model as equation 1.5 with the probability that a worker moves within the manager's network as the outcome.[71] I define a connected move based on whether the manager has ever worked (i) with the new manager the worker moves to and/or (ii) in the same sub-function and/or office as the job the worker moves to. Figure 1.24 shows no differential impact of gaining a high-flyer manager on connected moves, whether these are lateral or vertical moves (see Figure 1.44 for the reverse transition).

Third, I find that the higher career progression is unrelated to the workers' degree of homophily with the manager, such as sharing the same gender. Panel (b) in Figure 1.27 shows that the higher salary increases and transfer rates do not appear to differ by the degree of homophily with one's manager (in terms of gender). Therefore, a high-flyer manager has a positive effect on workers' careers irrespective of their gender. Again, my estimates do not identify any *differential* effect between high and low-flyers; it could still be that both types

[71] As the research design is based on the first manager change, I cannot look at pre-trends, i.e. I can only construct the indicator for manager's formal ties to destination for the post-transition period.

exhibit bias towards workers of the same gender. Namely, while effects resulting from social connections do not explain my results, they may well exist inside the organization.

Managers changing the jobs around the workers, instead of re-shuffling workers to existing jobs. The analysis conducted takes as fixed the nature of the jobs that workers can get allocated into. However, rather than shaping the matching of workers to jobs, high-flyer managers might change the jobs available to match them better to the skills of the existing workers. To check for this, I test if high-flyer managers are more likely to change the type of jobs by replacing "old" jobs with "new" ones. I define a job to be new if it is not appearing in the previous months within a given team. Correspondingly, I define a job to be old if it is no longer appearing in subsequent months within a given team. I also compute the share of managerial jobs (work-level 2+) within the same sub-function. Figure 1.25 shows that there are no differential effects of high-flyer managers on new job titles created and old job titles destroyed within a team, and the share of managerial jobs within a sub-function. Hence, I do not find evidence of high-flyer managers changing the type of jobs by replacing old jobs with new ones: they are re-shuffling workers to existing jobs instead of changing the jobs around the workers.

Congestion effects. Given the high-flyer managers have a higher chance to be further promoted to work-level 3 (Table 1.4), a concern could be that the career progression effects of the workers exposed to high-flyers are in part driven by a career spillover effect (Bianchi et al., 2022): high-flyers, by being promoted faster, leave room for a promotion for one of their subordinates. Three facts alleviate this concern. First, the asymmetric effects of the impact of losing a high-flyer manager represent evidence against this possibility. One would expect a negative impact on the probability of a vertical transfer for the workers experiencing the *HtoL* transition when compared to those with the *HtoH* transition. Second, I check directly whether the workers moving from a low- to a high-flyer have a higher chance of taking the exact position of their manager when compared to the workers moving to another low-flyer. The share of workers taking the place of their managers is actually 1ppt higher for the workers in the *LtoL* transition (8.9%) as opposed to the *LtoH* (7.9%).[72] Third, institutionally, managerial promotion decisions are taken at a more aggregate level than the team level, at the sub-function level, which represents the same unit within which work-level 2 managers typically rotate as part of the rotation policy. Hence, a faster promotion of a high-flyer manager from work-level 2 to work-level 3 would open up a managerial position for all workers within a sub-function,

[72] In unreported results, I can replicate the event-study plots of the effects of gaining a high-flyer manager when leaving out the workers who at one point are promoted to exactly the respective managerial positions. The results remain unchanged.

irrespective if they are in a team supervised by a low-flyer versus a team supervised by a high-flyer.

Talent hoarding. The empirical strategy is not designed to detect managers' talent hoarding behavior as I am comparing manager switchers that have the same outgoing manager but a different incoming manager. Hence, any effect due to a manager rotation is net out by design.[73] Yet, it is helpful to understand how to relate my results to potential talent hoarding behavior.

The interpretation crucially depends on the correlation between talent hoarding behavior and being a high-flyer manager. If there is no correlation between high-flyer status and talent hoarding, then hoarding behavior would be orthogonal to being a high-flyer and could not explain my results. This would indicate that talent hoarding may well exist in the organization, but cannot be the reason for my findings. This is in line with Haegele (2021) that shows that managers' talent hoarding behavior is not correlated with manager characteristics.

If there is a positive correlation - high-flyer managers are more likely to engage in talent hoarding - then the results found are a lower bound for the impacts of high-flyer managers on workers' careers.

If there is a negative correlation - high-flyer managers are less likely to engage in talent hoarding - then my results could be explained by the fact that high-flyer managers are less prone to engage in talent hoarding in comparison to low-flyer managers. This would be perfectly consistent with the interpretation put forward. However, to add some nuance, it is worth bearing in mind that there are several pieces of evidence implying that the absence of talent hoarding cannot fully explain my empirical results.

In particular, if talent hoarding were to largely explain the results, the high-performers at baseline should be the main or only drivers behind the positive effects of gaining a high-flyer manager on transfers and salary growth rates. Since the bad managers would be the talent hoarders, they would deny movements out of the team only for the high performers as they want to keep them in their team. Hence, moving to a high-flyer manager should increase the transfer and promotion rates exclusively for high-performers at baseline while the low performers would be unaffected. Instead, I find that also the weakest links (i.e. the lowest performers at baseline) have higher transfer rates when gaining a high-flyer manager and improve their career path (panel b in Figure 1.29). Besides, under the hypothesis of low-flyers engaging in

[73] See Friebel and Raith (2022) for theoretical work on how different organizational designs may change managers' incentives to train subordinates and accurately represent their abilities, and Haegele (2021) for an empirical application that uses manager rotations to analyze talent hoarding behavior. These papers build on the intuition that, when a manager learns that they will move to a new position on a different team, they no longer have the incentive to hoard workers on their current team; hence, for workers whose manager will soon rotate, this creates a temporary window of time during which they are not subject to talent hoarding.

talent hoarding, one would expect high-performers to have higher retention rates when gaining a high-flyer manager as they would quit the low-flyer managers that are obstructing their career development by hoarding them. However, I do not find differential effects on exit by baseline performance (panel a and panel b in Figure 1.29). Appendix Figure 1.66 also shows that the workers that transfer do not report different answers on the engagement survey. This rules out that the workers are changing jobs because of escaping a manager that is hoarding workers, rather than the proposed interpretation of workers finding a better match in terms of their skills in the organization. The asymmetries in the career trajectories when losing a high-flyer manager provide a further test against the talent hoarding interpretation.

Comparing manager rotations with workers' transfers. It is important to understand the difference between the managers' lateral rotations across teams and the workers' transfers. The firm adopts different personnel strategies at different levels of the hierarchy.

For entry-level workers in work-level 1, the firm's objective is to find the area inside the company where they can thrive. Workers are encouraged to actively think about their skills, interests, and their future goals and to keep a continual dialogue with their line managers on career development. In other words, *exploration* is deemed more important than *exploitation*.

For work-level 2 managers, the objective is to train them within a given area and understand who would be capable of progressing to work-level 3, hence they conduct lateral rotations across teams. Again, *exploration* is deemed more important than *exploitation* although the former is typically conducted within a given sub-function. As opposed to employees in work-level 1, the nature of the job does not change substantially in these rotations, what changes are the people the manager interacts with and the projects. The exploration that the firm cares about in this case is the one required to find the right work-level 3 managers, who start to have bigger responsibilities such as setting strategy and making budget allocation decisions.

For work-level 3 (around 2,200 employees in the cross-section), the objective is to exploit the knowledge they have accumulated in the specific area and avoid frequent rotations, or *exploitation* is deemed more important than *exploration*.

For work-level 4 and above, there are a number of different considerations that aim to strike a balance between exploitation and exploration/getting relevant work experience for the executive suite: some rotations, e.g. across countries, are encouraged although they typically last longer. This is a very selected pool of employees at the top echelons of the multinational (around 500 employees in the cross-section).

1.7 Conceptual framework

To explain the managers' effects on workers' careers, I set up a model linking managerial quality to worker performance through on-the-job talent discovery and learning by doing. Through the lenses of the model, I discuss how talent discovery can be empirically distinguished from teaching, the most plausible alternative channel, and what are the minimal conditions I need for the general framework to reconcile with my results.[74] The objective is not to develop a realistic model of the role of managers in internal labor markets but rather to elucidate some of the essential lessons from the empirical results.

The elemental economic problem that arises with on-the-job talent discovery has been well understood by economists at least since Johnson (1978) and Jovanovic (1979). The optimal solution to experimentation problems draws on the "bandit" literature, which shows how to account for the trade-off between output now and information that can help increase output in the future. There are also studies that combine experimentation in a labor market with multiple job types (MacDonald (1982); Miller (1984)). However, these papers abstract away from the role of individual managers in revealing workers' talents. In my framework, I introduce managers' heterogeneity in quality and examine their differential impact on workers within a simple setup in which production depends on performing a variety of tasks and workers differ in their task-specific human capital.

1.7.1 Model setup

Consider a firm that combines managers (b), workers (i) and occupations (o). Output in an occupation is produced by combining multiple tasks, e.g. negotiating, programming, and managing personnel (Autor et al. (2003); Gibbons and Waldman (2004); Lazear (2009); Gathmann and Schönberg (2010)). Workers differ in their task-specific human capital (i.e. workers have multidimensional skills).

Managers also differ in their task-specific human capital but, for simplicity and given the focus of this paper, I hone in on one overall human capital dimension for them, namely, managerial skill. In particular, let managerial skill take one of two types: high (H) and low (L) quality managers. The manager type categorization can be conceptualized in two complementary ways: good managers have a higher level of each skill and/or good managers have a higher level of all the skills related to managing subordinates, such as mentoring, teaching and motivating workers.

[74] Friebel and Raith (2022) highlights this dual role of managers in the development and allocation of human capital in firms: they train junior employees and acquire private information about workers that is needed to allocate them to the right positions.

The basic intuition can be developed with a one-period setup: managers are assigned to workers in a random fashion[75], observe worker productivity, and decide the job allocation of the worker. Throughout, the emphasis is on managers, and the workers are non-strategic players that follow the manager's decisions.[76]

1.7.2 Workers and task-specific human capital

Occupations (o) are bundles of tasks and differ in the importance of each task for production. For simplicity, let there be two tasks (j): A and S (e.g. analytical and social). Let β_o^A be the weight on the analytical task and β_o^S be the weight on the social task. The weights, β_o^j, indicate how important a particular task j is for a given occupation o.[77]

Workers have observed productivity in each task j, which is determined by a person's initial endowment m_i^j in each task ("talent"), the experience accumulated in task j until time t, E_{it}^j, and a noise term (ϵ_{iot}):

$$\begin{aligned} p_{iot}^j &= \underbrace{E_{it}^j}_{\text{experience}} + \eta_{iot}^j \\ \text{where} \quad \eta_{iot}^j &= \underbrace{m_i^j}_{\text{innate task talent}} + \underbrace{\epsilon_{iot}^j}_{\text{noise}} \end{aligned} \tag{1.8}$$

where t is time in the labor market, $m_i^j \sim N(\mu^j, \sigma^j)$ and $\epsilon_{iot}^j \sim N(0, \sigma_\epsilon^j)$. The noise or luck shocks, ϵ_{iot}^j, are uncorrelated across people, occupations, and tasks, and $\epsilon_{iot}^j m_i^j$.

There is learning-by-doing in each task, which depends on the task intensity on the job:

$$E_{it}^j = \sum_{o'} (\beta_{o'}^j) O_{io't} \tag{1.9}$$

where $O_{io't}$ is tenure in each prior occupation o'. For example, a worker accumulates more analytical skills if she works in an occupation in which analytical skills are very important (i.e., with a large β_o). In contrast, she will not learn anything in tasks that she does not use in her occupation.

[75] In the empirical strategy, I isolate exogenous assignments as part of the firm's routine of re-shuffling managers to teams to train and screen work-level 2 managers.

[76] While extreme, this is supported by the empirical setting: managers have to approve and sponsor worker applications for transfers and promotions.

[77] The weights allow for both horizontal (the ratio of the weights indicates the relative importance of each task) as well as vertical job differentiation (the level of the weights indicates the task intensity). For this reason, the weights are not constrained to be between 0 and 1 (and hence cannot be interpreted as the share of time a worker spends on average in a given task in occupation o). As an example, occupations in managerial positions would exhibit higher returns to the same tasks than the entry-level analogs, hence they would have higher weights for every task even though the ratios of the weights may be identical.

Hence, worker i's overall productivity (P) in log units is given by:

$$lnP_{iot} = \beta_o^A p_{iot}^A + \beta_o^S p_{iot}^S$$

$$\longrightarrow lnP_{iot} = \underbrace{(\beta_o^A E_{it}^A + \beta_o^S E_{it}^S)}_{\bar{E}_{iot}=\text{task-specific experience}} + \underbrace{(\beta_o^A m_i^A + \beta_o^S m_i^S)}_{\bar{m}_{io}=\text{task match}} + \underbrace{(\beta_o^A \epsilon_{iot}^A + \beta_o^S \epsilon_{iot}^S)}_{\bar{\epsilon}_{iot}=noise} \tag{1.10}$$

Note that learning by doing creates occupational persistence. As workers accumulate more and more task-specific experience as they age, a distant occupational switch tends to become increasingly costly.

1.7.3 Managers

Managers observe worker productivity and decide the next worker job allocation to maximize expected worker productivity.[78] Hence, the manager solves the following problem:

$$\max_{\beta_o} \sum_j \beta_o^j \mathbb{E}(p_{i,t+1}^j)$$

If full information on each worker were available, managers would assign workers to jobs based on comparative advantage. Without full information, managers choose the allocation that maximizes productivity in expectation. Expected productivity depends on expected task match, which is inferred from the productivity realization in each task j:

$$\mathbb{E}(m_{it}^j) = \hat{m}_{iot}^j = p_{iot}^j - E_{it}^j = m_i^j + \epsilon_{iot}^j$$

I allow good and bad managers to differ in two fundamental ways: in terms of solving the job assignment problem based on the expected task talents (*matching channel*); and in terms of influencing the speed of workers' learning-by-doing (*teaching channel*).

First, the *matching channel*: while bad managers infer workers' innate talents based on the productivity realization, $\hat{m}_{iot}^j = m_i^j + \epsilon_{iot}^j$, good managers receive a private signal that enables them to fully discover the workers' talents, m_i^j (one-shot learning process). Managers use this information to potentially re-optimize the job allocation decision. Given that the good manager has fully revealed the worker's innate talents, future worker productivity is higher on average as the workers locate better matches.

[78] In my setting, managers' input is key in any transfers and promotions decisions and a worker cannot apply to another position inside the company without the manager's approval. In this framework, I am not considering the manager's incentives. This is supported by the empirical strategy that compares outcomes between different types of managers, netting out common managerial behaviors due to the firm's policies.

Second, the *teaching channel*: good managers increase the speed of workers' learning-by-doing. Experience on the job depends on the manager's quality as follows:

$$E^j_{it} = \sum_{o'} (\beta^j_{o'} (I_{b_{o'}=Good})\tau) O_{io't}$$

where $I_{b_{o'}=Good}$ is the indicator function for good manager and $\tau > 1$. After one period of working under a good manager, a worker has accumulated more on-the-job experience compared to working under a bad manager. There can be different reasons why good managers may increase workers' on-the-job experience such as teaching and training activities or motivating workers to exert higher effort.

1.7.4 Model predictions

I now illustrate how the productivity and transfer dynamics depend on the manager of the worker. Let there be two jobs: one mostly analytical ($\beta^A = 1 - \delta; \beta^S = \delta$) and one mostly social ($\beta^A = \delta$; $\beta^S = 1 - \delta$), with $\delta \to 0$ (δ is infinitesimally small). Hence, the manager can observe the task-specific productivity for each task (as $\delta > 0$) but only one task basically matters for each job (given that $\delta \to 0$). The worker starts with no experience in either the analytical or social job; for simplicity and without loss of generality, the initial job allocation is assumed to be orthogonal to the worker's innate talents. Let the worker have higher analytical skills $m^A > m^S$, thus output would be maximized by allocating the worker to the analytical job.

The dynamics will depend on the initial job allocation. Table 1.21 shows how the expected worker productivity computed by the manager changes depending on the manager type and the job allocation. As a reminder, a good manager perfectly observes a worker's innate talents. Using Table 1.21, I can derive the following predictions:

Prediction 1, good manager. *A good manager moves a worker from job o' to job o if:*

$$\underbrace{(\bar{m}_{iot} - \bar{m}_{io't})}_{\Delta\bar{m}_{iot}=\text{ gain in task match}} > \underbrace{\bar{E}_{io't} - \bar{E}_{iot}}_{-\Delta\bar{E}_{iot}=\text{ potential loss in task-specific experience}} \tag{1.11}$$

that is if the allocation gain outweighs the teaching loss or the matching channel is more important than the teaching channel.

Hence, given the example above, a good manager moves the worker from the social to the analytical job if:

$$\underbrace{m^A - m^S}_{\text{gain in task match}} > \underbrace{\tau}_{\text{loss in task-specific experience}}$$

On the other hand, a good manager never moves the worker from the analytical to the social job. If the worker starts in the analytical job, she is well-matched according to her talents and the teaching channel via learning-by-doing reinforces the gains of an initial good allocation.

Prediction 2, bad manager. *A bad manager moves a worker from job o' to job o if:*

$$\underbrace{(\hat{\bar{m}}_{iot} - \hat{\bar{m}}_{io't})}_{\Delta\hat{\bar{m}}_{iot} = \text{ gain in expected task match}} > \underbrace{\bar{E}_{io't} - \bar{E}_{iot}}_{-\Delta\bar{E}_{iot} = \text{ potential loss in task-specific experience}} \tag{1.12}$$

that is, a worker is assigned to a different job if the improvement in the expected task match exceeds the potential loss in task-specific experience.

Hence, given the example above, a bad manager moves the worker from the social to the analytical job if:

$$(m^A + \epsilon_1^A) - (m^S + \epsilon_1^S) > 1 \Rightarrow (\epsilon_1^A - \epsilon_1^S) > 1 - (m^A - m^S)$$

that is, the probability of a bad manager moving the worker is given by $1 - \Phi\left(\frac{1-(m^A-m^S)}{\sigma_\epsilon^{2A}+\sigma_\epsilon^{2S}}\right) = \Phi\left(\frac{(m^A-m^S)-1}{\sigma_\epsilon^{2A}+\sigma_\epsilon^{2S}}\right)$, by symmetry of the standard normal distribution and if Prediction 1 holds $\left(m^A - m^S > \tau\right)$.

Similarly, a bad manager moves the worker from the analytical to the social job if:

$$(m^S + \epsilon_1^S) - (m^A + \epsilon_1^A) > 1$$

that is, the probability of a bad manager moving the worker is $1 - \Phi\left(\frac{(m^A-m^S)+1}{\sigma_\epsilon^{2A}+\sigma_\epsilon^{2S}}\right)$. The two moving probabilities do not sum to one given the experience term that accumulates via learning by doing.

1.7.5 Discussion: mapping predictions to the empirical results

By combining Prediction 1 and 2, I can now discuss how moves and productivity depend on the worker history in terms of manager types, thus mapping the empirical research design. In particular, I am interested in the predictions around gaining a good manager (moving from a low- to a high-flyer compared to moving from a low- to another low-flyer manager) and losing a good manager (moving from a high- to a low-flyer compared to moving from a high- to another high-flyer manager). The results indicate that gaining a good manager has a positive impact on both transfers and productivity while losing a good manager has no impact.

The empirical findings are consistent with the matching channel being more important than the teaching channel, in terms of what differentiates good managers from the rest, and the model makes clear that the effects of a good versus a bad manager depend on the worker's history in the following way. If the previous manager was bad, there is a non-zero probability of job misallocation and hence the good manager can increase the worker productivity by changing her job allocation. If the matching channel is more important than the teaching channel, then a good manager leads to higher transfers and productivity compared to a bad manager. On the other hand, if the previous manager was good, the probability of job misallocation is zero, as workers have already been assigned to jobs according to their talents, and hence another good manager is not having an additional impact if there is no difference among manager types in teaching efficacy or if there are decreasing returns to learning-by-doing.[79] The bad manager is also not going to move a worker given she observes that the worker had a good manager before.

The model illustrates how, through talent discovery, there are dynamic benefits of having had a good manager at least once during a worker's career. The asymmetries in the empirical results are consistent with this intuition (see sub-section 1.6.1). It also follows that the predictions should be stronger when the worker's initial labor market experience is low (low $\bar{E}_{iot}$, for e.g. younger workers)[80], as found in the heterogeneity analysis also discussed in sub-section 1.6.1.

Appendix 1.14 includes additional details and specifies more precise conditions on the parameters for these predictions.

1.8 Conclusion

Managers are at the heart of organizations, within which they determine the allocation of resources, and thus fundamental in the theory of the firm (Coase (1937); Chandler (1977)). Their importance can also be seen in the latest empirical trends: globally, the managers' share of wages is 38% (ILO, 2019). And yet, empirical evidence studying the long-term impact of individual managers on workers' careers and its link to firm-level outcomes remains sparse. I open the "black box" of the firm by collecting novel personnel records from a large consumer goods multinational and provide evidence that the ability of managers to match diversely skilled workers to specialized jobs inside the firm has large and persistent effects on the worker performance and career path, as well as on the productivity of an establishment as a whole.

The impacts of a worker's exposure to a good manager extend far beyond the period circumscribed by the particular manager-worker spell. In fact, it may often be through the future career development of their workers that managers' greatest influence on firm productivity occurs. Such gains are out of a more productive allocation of workers and occur potentially at

[79]For the predictions from the simple model to match exactly with the empirical results, either the teaching has to be the same between a good and a bad manager or there have to be decreasing returns to experience, which are both plausible.

[80]When experience is low, there is more scope for gains out of a job re-allocation as captured by equation 1.12.

zero cost, as they do not require any firing, hiring, or training of workers.

In studying the internal labor market of a multinational firm, I extend the grasp of economic analysis to questions of importance to today's large companies. This is particularly relevant when considering that, across the OECD countries, large firms with over 250 workers represent only 1% of enterprises but account for a staggering 40% of manufacturing employment.[81] Modern business enterprises feature rich and complex internal labor markets characterized by a multiplicity of horizontally differentiated jobs as well as vertical layers. Within these, firms rely on managers to determine the allocation of workers to jobs and to steer workers' careers so that they can reach their potential in the organization (Drucker (2001); Conaty and Charan (2010)).[82] While the results pertain to only one firm and the magnitude of the effects may vary in other contexts, the mechanism of managers harnessing workers' unique skills by directing them to their most suitable career path is of general application.

Considering managerial training and management practices, my results highlight the allocation of workers to jobs as an important margin for improving performance. The ability to create efficient worker-job matches is particularly valuable at times when technological innovation such as digitalization and artificial intelligence, and disruptions such as pandemics or climate change, force widespread firm restructuring and require the reallocation of existing workers to new jobs or their replacement with workers featuring new skills. Moreover, my results imply that the most successful managers (as identified by the firm) are able to extract more value from the same managerial practices set by firm-wide policies, indicating that the effectiveness of managerial practices also depends on the managers' ability to use them.

Methodologically, instead of using surveys regarding the way managers run their firms' operations, I analyze rich administrative firm data, unpacking the managers' impacts by looking at outcomes from *within* the firm. The data does not shed light on the precise skills needed for managers to enable the discovery of workers' unique aptitudes and whether managers can get trained in these or whether they are innate. Designing effective training initiatives to test this as well as understanding if predictions by artificial intelligence can substitute for or complement human skills are fascinating questions for future research.

[81] Based on OECD Structural and Demographic Business Statistics. Autor et al. (2020) documents extensively how large firms are getting bigger over the last five decades across high-income countries. For example, the share of U.S. employment in firms with more than 5,000 employees rose from 28% in 1987 to 34% in 2016.

[82] Organizations such as General Electric, Procter and Gamble, LG, and Novartis have been heavily investing in building effective people management strategies to develop and allocate employees to the positions they are best suited for. Increasingly, the responsibility for talent management is shifting from HR to frontline managers (Whittaker and Marchington (2003); Perry and Kulik (2008); Cappelli (2013)).

1.9 References

Acemoglu, Daron and David Autor, "Skills, tasks and technologies: Implications for employment and earnings," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1043–1171.

Adhvaryu, Achyuta, Anant Nyshadham, and Jorge Tamayo, "Managerial Quality and Productivity Dynamics," 2022.

_ , Namrata Kala, and Anant Nyshadham, "Management and shocks to worker productivity," 2022.

Ashraf, Nava, Oriana Bandiera, Virginia Minni, and Víctor Quintas-Martínez, "Gender Roles and the Misallocation of Labour Across Countries," 2022.

Autor, David, "The "task approach" to labor markets: an overview," *Journal for Labour Market Research*, 2013, *46* (3), 185–199.

_ , David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen, "The fall of the labor share and the rise of superstar firms," *The Quarterly Journal of Economics*, 2020, *135* (2), 645–709.

Autor, David H, Frank Levy, and Richard J Murnane, "The skill content of recent technological change: An empirical exploration," *The Quarterly journal of economics*, 2003, *118* (4), 1279–1333.

Bagger, Jesper, Fran Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin, "Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics," *American Economic Review*, 2014, *104* (6), 1551–96.

Baker, George and Bengt Holmstrom, "Internal labor markets: Too many theories, too few facts," *The American Economic Review*, 1995, *85* (2), 255–259.

_ , Michael Gibbs, and Bengt Holmstrom, "The internal economics of the firm: Evidence from personnel data," *The Quarterly Journal of Economics*, 1994, *109* (4), 881–919.

_ , _ , and _ , "The wage policy of a firm," *The Quarterly Journal of Economics*, 1994, *109* (4), 921–955.

Bandiera, Oriana, Andrea Prat, Stephen Hansen, and Raffaella Sadun, "CEO behavior and firm performance," *Journal of Political Economy*, 2020, *128* (4), 1325–1369.

— , **Iwan Barankay, and Imran Rasul**, "Incentives for managers and inequality among workers: Evidence from a firm-level experiment," *The Quarterly Journal of Economics*, 2007, *122* (2), 729–773.

— , — , **and** — , "Team incentives: Evidence from a firm level experiment," *Journal of the European Economic Association*, 2013, *11* (5), 1079–1114.

Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta, "Measuring and analyzing cross-country differences in firm dynamics," in "Producer dynamics: New evidence from micro data," University of Chicago Press, 2009, pp. 15–76.

Benson, Alan, Danielle Li, and Kelly Shue, "Potential" and the gender promotion gap," Technical Report, working paper 2021.

Bertrand, Marianne and Antoinette Schoar, "Managing with style: The effect of managers on firm policies," *The Quarterly journal of economics*, 2003, *118* (4), 1169–1208.

— , **Robin Burgess, Arunish Chawla, and Guo Xu**, "The glittering prizes: Career incentives and bureaucrat performance," *The Review of Economic Studies*, 2020, *87* (2), 626–655.

Bhagat, Sanjai, Andrei Shleifer, Robert W Vishny, Gregg Jarrel, and Lawrence Summers, "Hostile takeovers in the 1980s: The return to corporate specialization," *Brookings Papers on Economic Activity. Microeconomics*, 1990, *1990*, 1–84.

Bianchi, Nicola, Giulia Bovini, Jin Li, Matteo Paradisi, and Michael L Powell, "Career spillovers in internal labor markets," 2022.

Bloom, Nicholas and John Van Reenen, "Human resource management and productivity," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1697–1767.

Cappelli, Peter, "HR for neophytes," 2013.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

Chade, Hector, Jan Eeckhout, and Lones Smith, "Sorting through search and matching models in economics," *Journal of Economic Literature*, 2017, *55* (2), 493–544.

Chaisemartin, Clément De and Xavier d'Haultfoeuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–96.

Chandler, Alfred D., "The visible hand: the managerial revolution in American business," *Cambridge, Massachusetts*, 1977.

Clifton, Jim and James K Harter, *It's the Manager: Gallup finds the quality of managers and team leaders is the single biggest factor in your organization's long-term success*, Gallup Press Washington, DC, USA:, 2019.

Coase, Ronald H., "The Nature of the Firm," *Economica*, 1937, *4* (16), 386–405.

Conaty, Bill and Ram Charan, *The talent masters: Why smart leaders put people before numbers*, Crown Business, 2010.

Correia, Sergio, Paulo Guimarães, and Tom Zylkin, "Fast Poisson estimation with high-dimensional fixed effects," *The Stata Journal*, 2020, *20* (1), 95–115.

Cortes, Guido Matias and Giovanni Gallipoli, "The costs of occupational mobility: An aggregate analysis," *Journal of the European Economic Association*, 2018, *16* (2), 275–315.

Cullen, Zoë B and Ricardo Perez-Truglia, "The old boys' club: Schmoozing and the gender gap," Technical Report, National Bureau of Economic Research 2019.

Davis, Steven J, John Haltiwanger, Kyle Handley, Ron Jarmin, Josh Lerner, and Javier Miranda, "Private equity, jobs, and productivity," *American Economic Review*, 2014, *104* (12), 3956–90.

Doeringer, Peter B and Michael J Piore, "Internal labor markets and manpower analysis, Lexington," *Ma: DC Heath*, 1971.

Drucker, Peter F, *Management Challenges for the 21st Century*, HarperBusiness edition, 2001.

Fenizia, Alessandra, "Managers and productivity in the public sector," *Econometrica*, 2022, *90* (3), 1063–1084.

Foster, Lucia, John C Haltiwanger, and Cornell John Krizan, "Aggregate productivity growth: Lessons from microeconomic evidence," in "New developments in productivity analysis," University of Chicago Press, 2001, pp. 303–372.

Frederiksen, Anders, Lisa B Kahn, and Fabian Lange, "Supervisors and performance management systems," *Journal of Political Economy*, 2020, *128* (6), 2123–2187.

Friebel, Guido and Michael Raith, "Talent Management: The Role of Bosses," 2022.

Gathmann, Christina and Uta Schönberg, "How general is human capital? A task-based approach," *Journal of Labor Economics*, 2010, *28* (1), 1–49.

Gelbach, Jonah B, "When do covariates matter? And which ones, and how much?," *Journal of Labor Economics*, 2016, *34* (2), 509–543.

Gibbons, Robert and Michael Waldman, "A theory of wage and promotion dynamics inside firms," *The Quarterly Journal of Economics*, 1999, *114* (4), 1321–1358.

_ and _ , "Task-specific human capital," *American Economic Review*, 2004, *94* (2), 203–207.

_ and Rebecca Henderson, *What do managers do?: Exploring persistent performance differences among seemingly similar enterprises*, in the The Handbook of Organizational Economics edited by Gibbons, R., Roberts, J., 2012.

Guvenen, Fatih, Burhan Kuruscu, Satoshi Tanaka, and David Wiczer, "Multidimensional skill mismatch," *American Economic Journal: Macroeconomics*, 2020, *12* (1), 210–44.

Haegele, Ingrid, "Talent Hoarding in Organizations," Technical Report, Working paper 2021.

Heckman, James J and Rodrigo Pinto, "Econometric mediation analyses: Identifying the sources of treatment effects from experimentally estimated production technologies with unmeasured and mismeasured inputs," *Econometric reviews*, 2015, *34* (1-2), 6–31.

Hoffman, Mitchell and Steven Tadelis, "People management skills, employee attrition, and manager rewards: An empirical analysis," *Journal of Political Economy*, 2021, *129* (1), 243–285.

Holmström, Bengt, "Moral hazard and observability," *The Bell journal of economics*, 1979, pp. 74–91.

Hsieh, Chang-Tai and Peter J Klenow, "Misallocation and manufacturing TFP in China and India," *The Quarterly journal of economics*, 2009, *124* (4), 1403–1448.

_ , Erik Hurst, Charles I Jones, and Peter J Klenow, "The allocation of talent and us economic growth," *Econometrica*, 2019, *87* (5), 1439–1474.

Imai, Kosuke, Luke Keele, and Dustin Tingley, "A general approach to causal mediation analysis.," *Psychological methods*, 2010, *15* (4), 309.

_ , _ , and Teppei Yamamoto, "Identification, inference and sensitivity analysis for causal mediation effects," *Statistical science*, 2010, *25* (1), 51–71.

Jaffe, Adam, "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value," *American Economic Review*, 1986, *76* (5), 984–1001.

Jaffe, Adam B, "Real effects of academic research," *The American economic review*, 1989, pp. 957–970.

Johnson, William R, "A theory of job shopping," *The Quarterly Journal of Economics*, 1978, pp. 261–278.

Jovanovic, Boyan, "Job matching and the theory of turnover," *Journal of political economy*, 1979, *87* (5, Part 1), 972–990.

Kahn, Lisa B and Fabian Lange, "Employer learning, productivity, and the earnings distribution: Evidence from performance measures," *The Review of Economic Studies*, 2014, *81* (4), 1575–1613.

Lazear, Edward P, "Performance pay and productivity," *American Economic Review*, 2000, *90* (5), 1346–1361.

_ , "Firm-specific human capital: A skill-weights approach," *Journal of political economy*, 2009, *117* (5), 914–940.

_ **and Sherwin Rosen**, "Rank-order tournaments as optimum labor contracts," *Journal of political Economy*, 1981, *89* (5), 841–864.

_ **, Kathryn L Shaw, and Christopher T Stanton**, "The value of bosses," *Journal of Labor Economics*, 2015, *33* (4), 823–861.

Lise, Jeremy and Fabien Postel-Vinay, "Multidimensional skills, sorting, and human capital accumulation," *American Economic Review*, 2020, *110* (8), 2328–76.

MacDonald, Glenn M, "A market equilibrium theory of job assignment and sequential accumulation of information," *The American Economic Review*, 1982, *72* (5), 1038–1055.

Metcalfe, Robert D, Alexandre B Sollaci, and Chad Syverson, "PRELIMINARY: Managers and Productivity in Retail," 2022.

Miller, Robert A, "Job matching and occupational choice," *Journal of Political economy*, 1984, *92* (6), 1086–1120.

Morris, Carl N, "Parametric empirical Bayes inference: theory and applications," *Journal of the American statistical Association*, 1983, *78* (381), 47–55.

Moscarini, Giuseppe, "Job matching and the wage distribution," *Econometrica*, 2005, *73* (2), 481–516.

Pastorino, Elena, "Careers in firms: the role of learning and human capital," Technical Report, Technical report, mimeo 2019.

Perry, Elissa L and Carol T Kulik, "The devolution of HR to the line: Implications for perceptions of people management effectiveness," *The International Journal of Human Resource Management,*, 2008, *19* (2), 262–273.

Rosen, Sherwin, "The theory of equalizing differences," *Handbook of labor economics*, 1986, *1*, 641–692.

Smith, Adam, *The wealth of nations [1776]*, Vol. 11937, na, 1937.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2020.

van Houten, Gijs and Giovanni Russo, "European Company Survey 2019-Workplace practices unlocking employee potential," 2020.

Waldman, Michael, "Job assignments, signalling, and efficiency," *The RAND Journal of Economics*, 1984, *15* (2), 255–267.

Whittaker, Susan and Mick Marchington, "Devolving HR responsibility to the line: threat, opportunity or partnership?," *Employee Relations*, 2003.

1.10 Figures

Figure 1.1: Distribution of work-levels by tenure (cumulative)

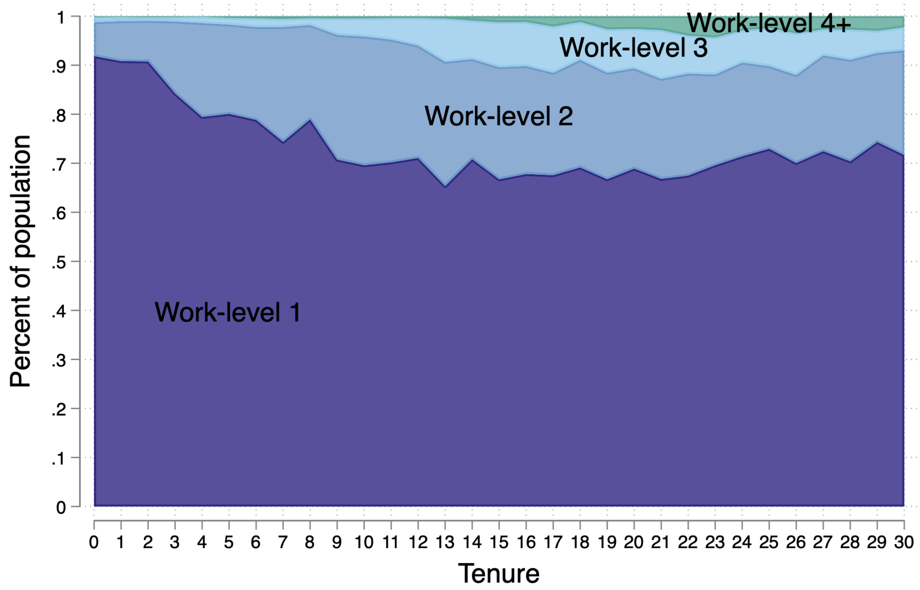

Notes. This figure shows the distribution of work-levels at different tenure years.

Figure 1.2: Age at promotion: work-level 2 managers

Notes. This figure shows the distribution of the age at promotion to work-level 2.

Figure 1.3: Manager value added in workers' future pay and high-flyer status

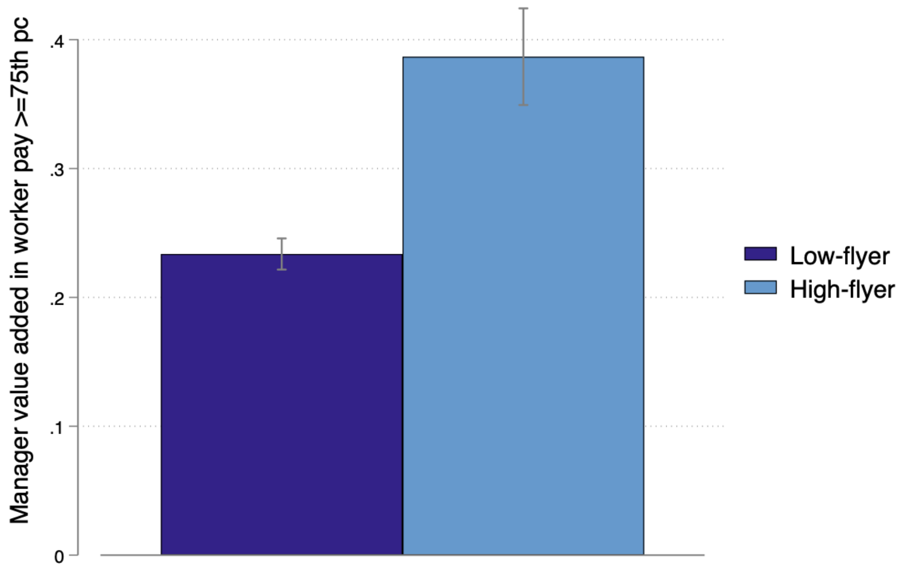

Notes. This figure shows the correlation between high-flyer status and the managers' value added in workers' pay being above the 75th percentile. I take worker pay in logs with a five year gap with respect to the manager exposure. I perform an Empirical Bayes shrinkage procedure for the fixed effects estimates to take into account of upward bias in the variance due to sampling noise (Morris, 1983).

Figure 1.4: Effects of gaining a high-flyer manager on lateral transfers, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s})$

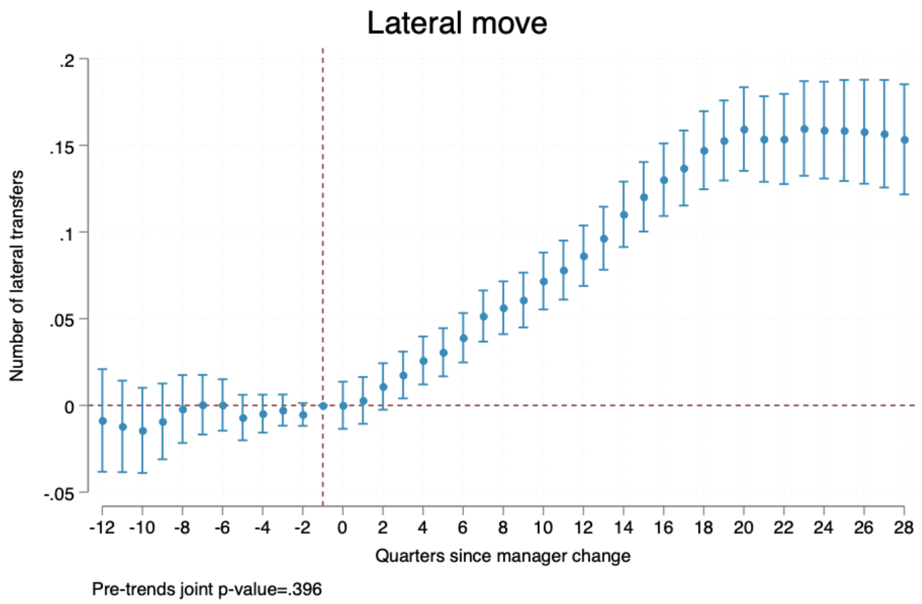

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The outcome variable is number of lateral transfers.

Figure 1.5: Effects of gaining a high-flyer manager on cross-functional transfers, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s})$

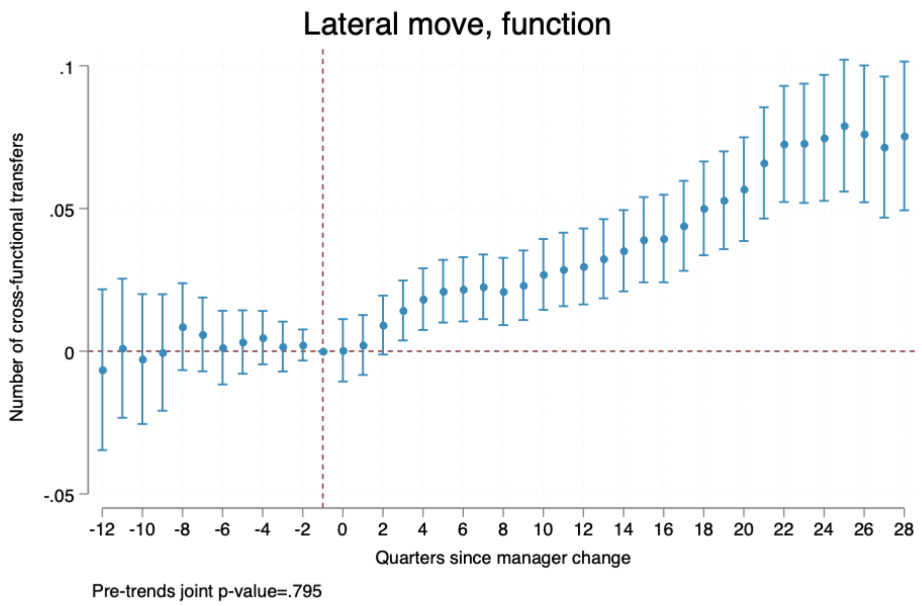

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. A cross-functional transfer is defined as a transfer across the 14 functions at the firm, e.g. from Finance to R&D.

Figure 1.6: Effects of gaining a high-flyer manager on task distance in transfers (O*NET), $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s})$

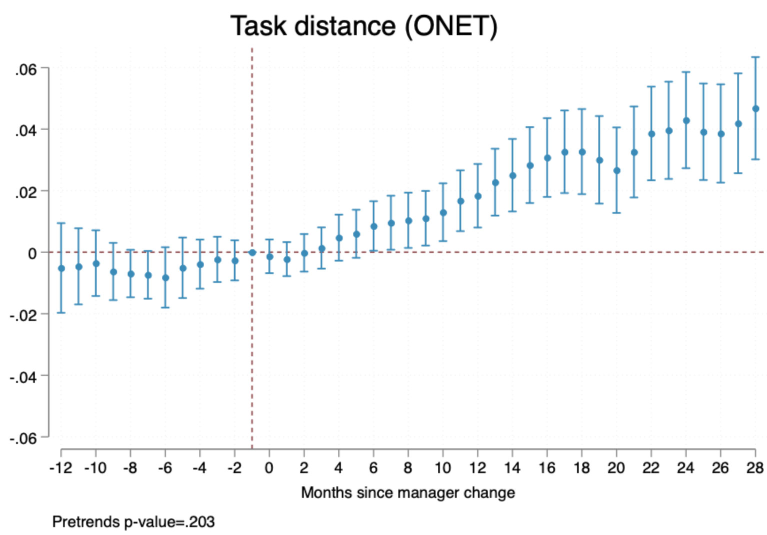

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. Task distance across jobs is constructed by matching the firm's job titles with O*NET data.

Figure 1.7: Effects of gaining a high-flyer manager on work-level promotions, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s})$

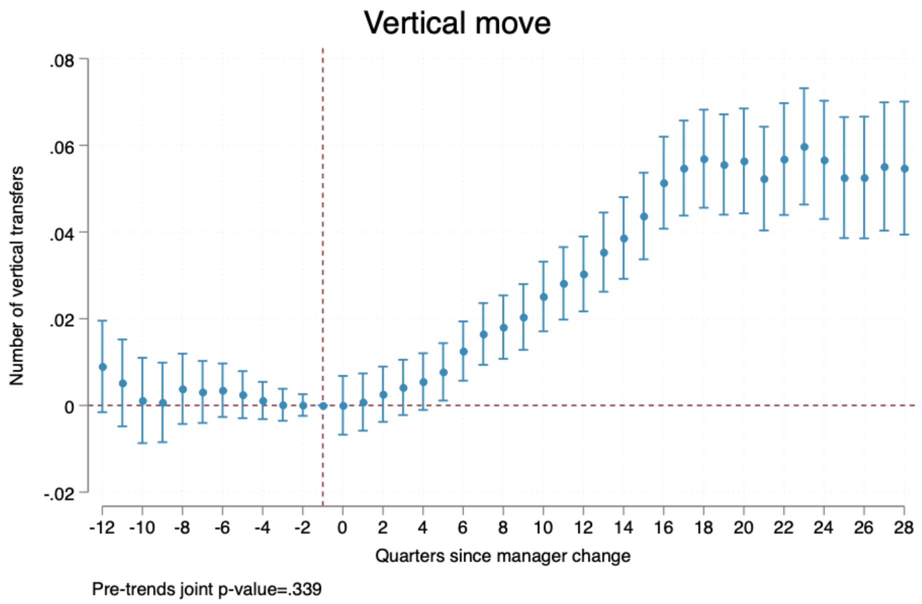

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The outcome variable is number of work-level promotions.

Figure 1.8: Effects of gaining a high-flyer manager on exit from the firm, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s})$

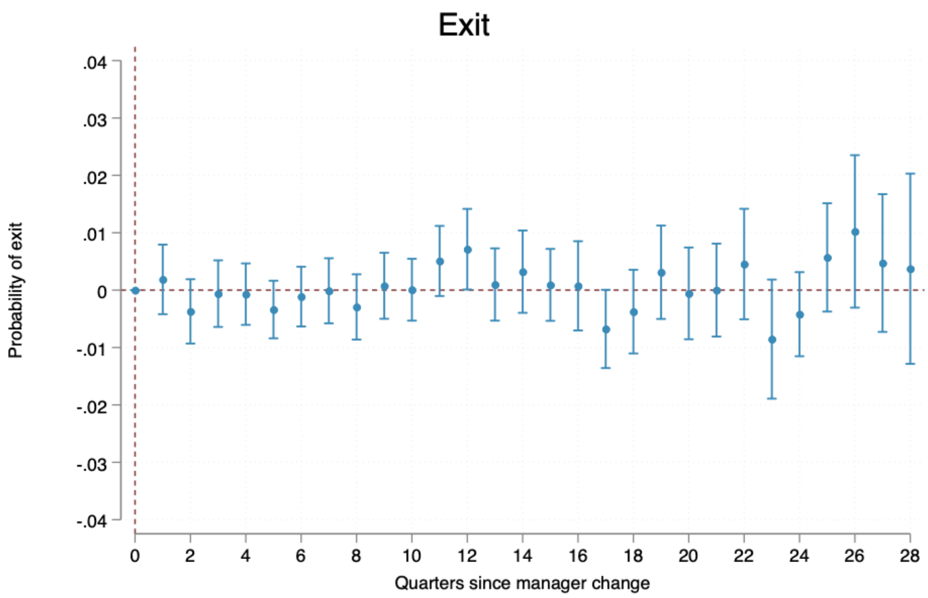

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The outcome variable is worker exit from the firm.

Figure 1.9: Effects of gaining a high-flyer manager on quits and layoffs, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s})$

(a) Quits **(b)** Layoffs

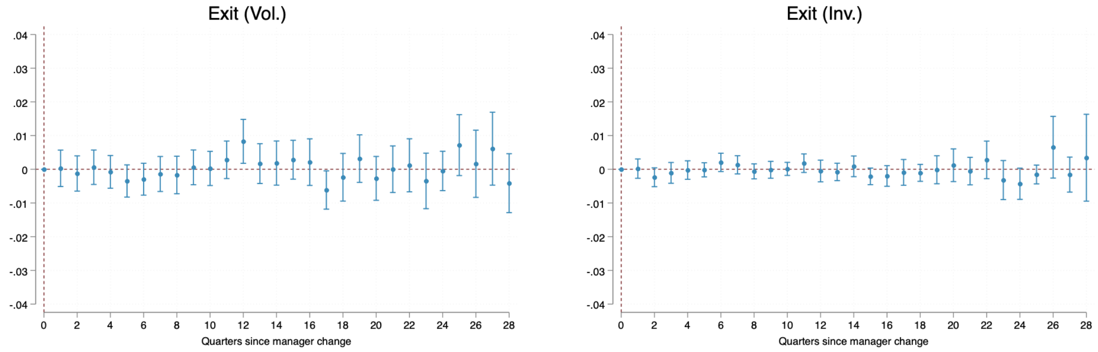

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The outcome variable is worker voluntary and involuntary exit from the firm.

Figure 1.10: Effects of gaining a high-flyer manager on salary grade increases, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s})$

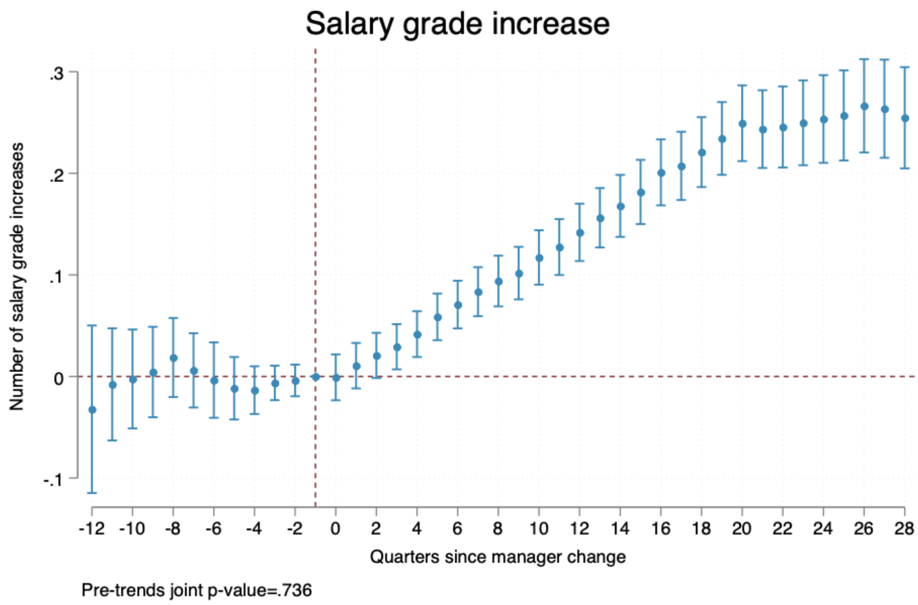

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The outcome variable is number of salary grade increases.

Figure 1.11: Effects of gaining a high-flyer manager on pay and bonus, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s})$

(a) Pay + bonus (logs) **(b)** Bonus (logs)

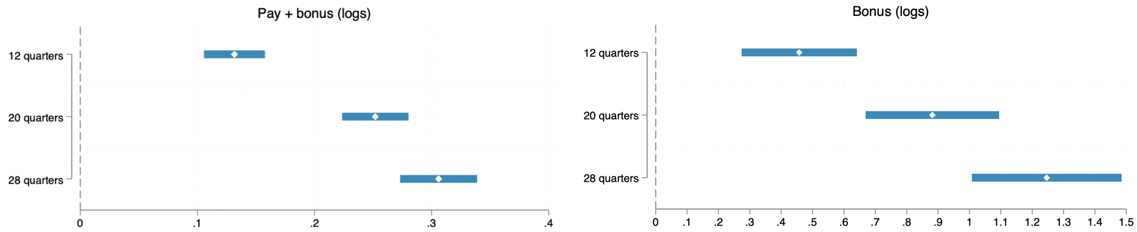

Notes. An observation is a team-year-month. Aggregating the monthly coefficients to the quarterly level. Reporting the estimates at 12, 20 and 28 quarters after the manager transition. 90% confidence intervals used and standard errors clustered by manager.

Figure 1.12: Effects of gaining a high-flyer manager on pay dispersion within the team, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s})$

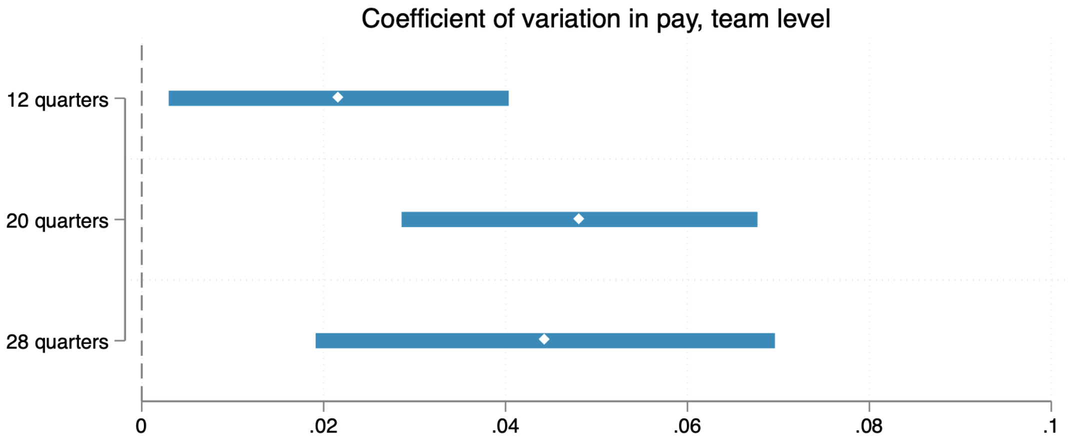

Notes. An observation is a team-year-month. Aggregating the monthly coefficients to the quarterly level. Reporting the estimates at 12, 20 and 28 quarters after the manager transition. 90% confidence intervals used and standard errors clustered by manager.

Figure 1.13: Gaining a high-flyer manager: team performance

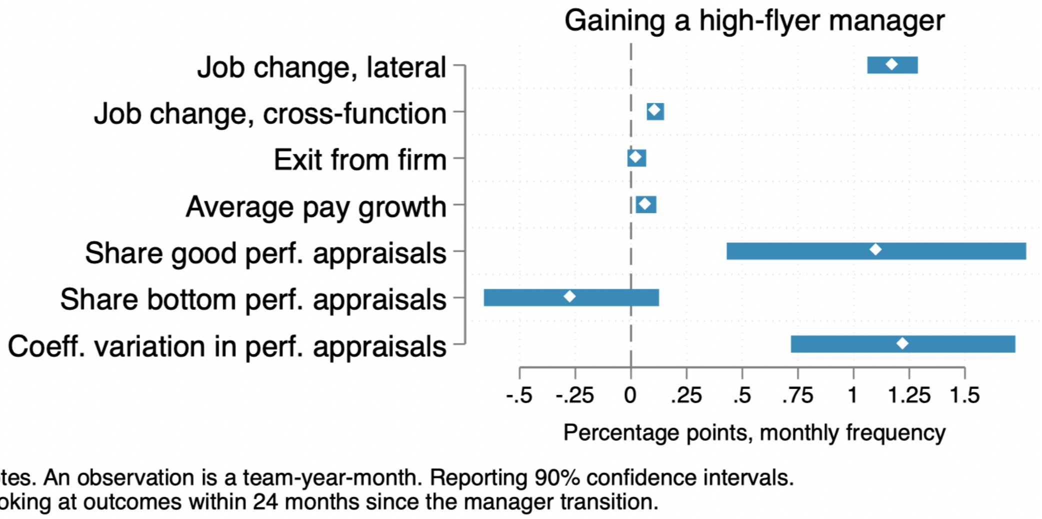

Notes. An observation is a team-year-month. 90% confidence intervals used and standard errors clustered by manager. Outcome mean, low-flyer: job change, lateral = 6ppt; job change, cross-functional = 0.4ppt; exit=1ppt; pay growth= 0.24ppt; good performance appraisals = 34ppt; bottom performance appraisals=11ppt.

Figure 1.14: Gaining a high-flyer manager: decomposing lateral moves

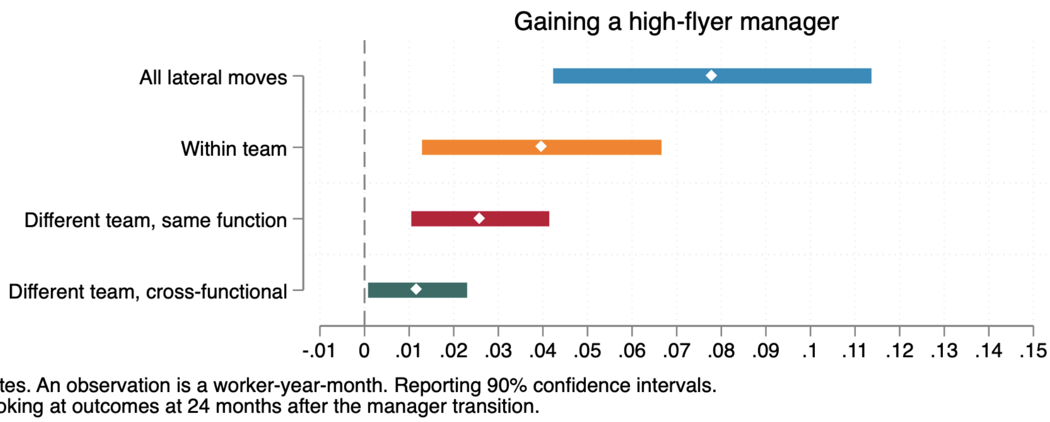

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Coefficients are estimated from a regression as in equation 1.2 and the figure reports the coefficient at the 8th quarter since the manager transition (since the average duration of a manager's assignment is two years).

Figure 1.15: High-flyer managers and objective worker performance

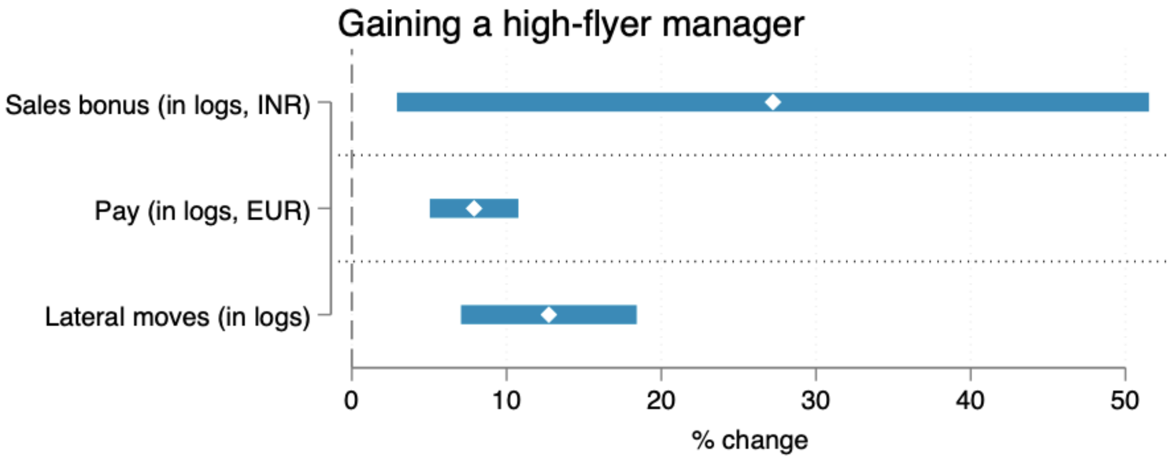

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Estimates obtained by running the model in equation 1.4. Controls include: worker FE and year-month FE. The outcome mean under a low-flyer manager is: Sales bonus (INR) 9,800; Pay (EUR) 10,600; Lateral Moves = 0.67.

Figure 1.16: High-flyer managers, lateral moves, and objective worker performance

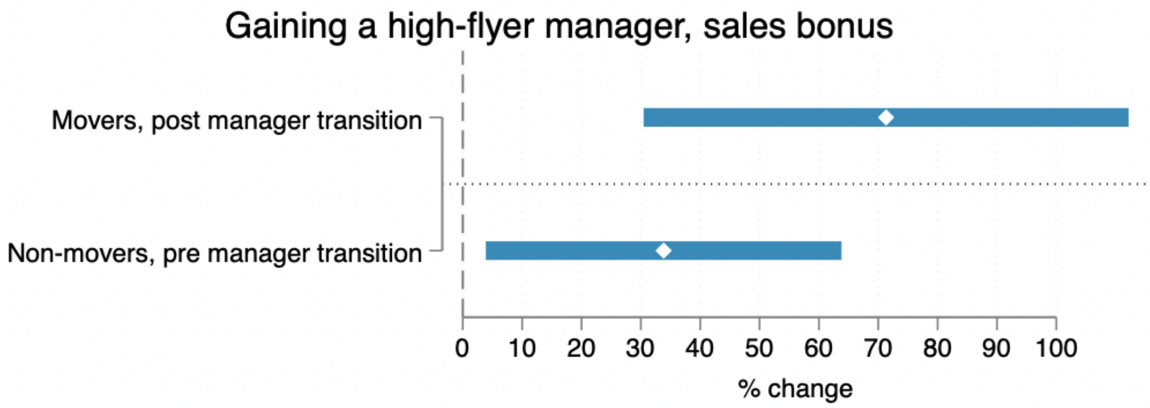

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Estimates obtained by running the model in equation 1.4. Controls include: worker FE (only for first row) and year-month FE. The first row looks at the the impact of gaining a high-flyer manager on sales bonus for workers that make at least one lateral move after the manager transition (up to five years after). The second row looks at the impact of gaining a high-flyer manager on sales bonus before a manager transition for the workers that do not make a job move after the manager transition.

Figure 1.17: Factory productivity (output per worker) and past exposure to high-flyer managers

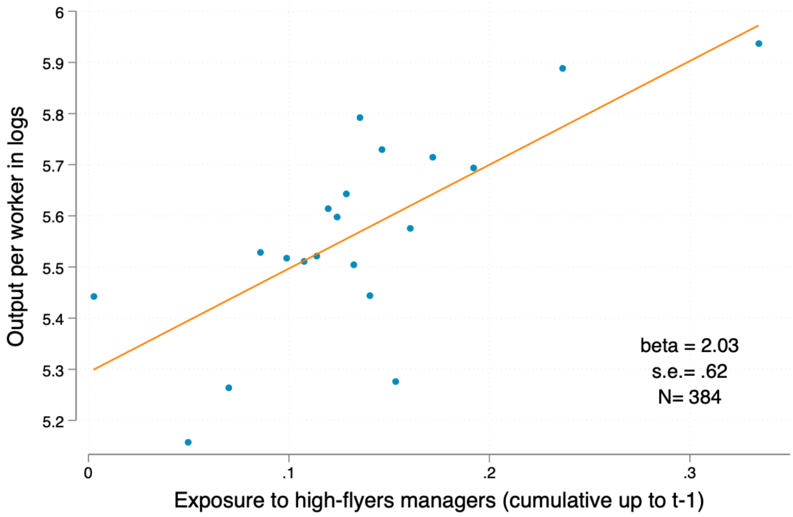

Notes. An observation is a factory-year. Standard errors clustered by factory-year. The y-axis is output per worker in logs (tons per worker) and the x-axis is the workers' cumulative exposure to high-flyers up to the year before. Controls include: country, product category and year fixed effects, share of managers, number of blue-collar and white-collar workers.

Figure 1.18: Factory costs per ton and past exposure to high-flyer managers

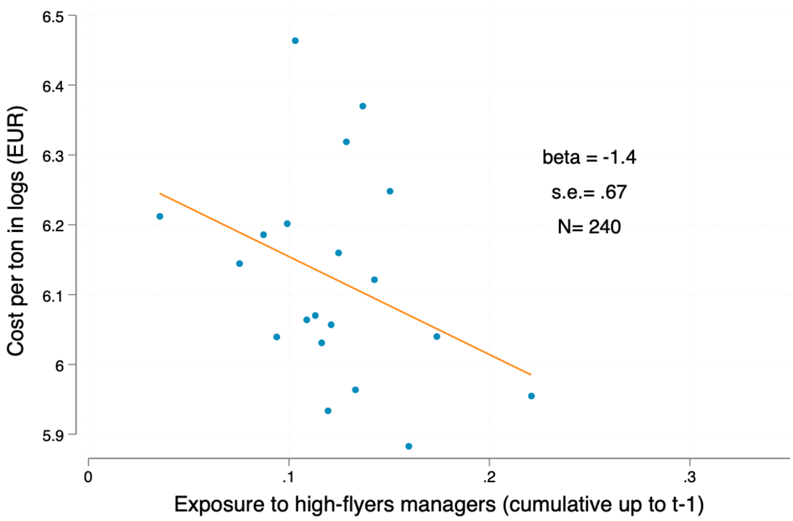

Notes. An observation is a factory-year. Because of changing reporting requirements, the costs per ton data could only be shared for the main product category (there are three product categories in total). Standard errors clustered by factory-year. The y-axis is operational costs per ton (EUR) in logs and the x-axis is the workers' cumulative exposure to high-flyers up to the year before. Controls include: country, and year fixed effects, share of managers, number of blue-collar and white-collar workers.

Figure 1.19: Factory productivity (output per worker) and workers' past lateral moves

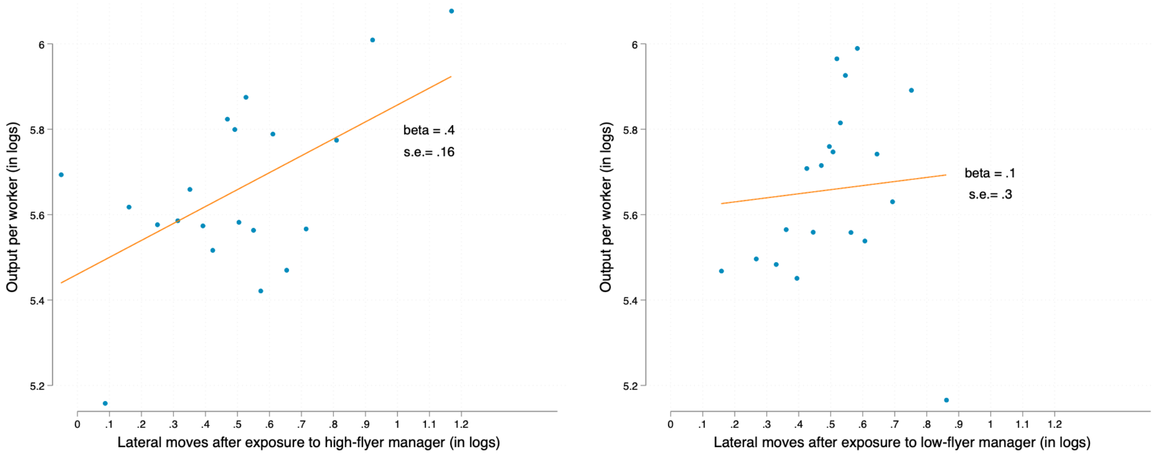

Notes. An observation is a factory-year. Standard errors clustered by factory-year. The y-axis is output per worker in logs (tons per worker) and the x-axis is the number of workers' lateral moves in logs: a) the left panel only considers the lateral moves that occur up to five years after being exposed to a high-flyer manager; b) the right panel only considers the lateral moves that occur up to five years after being exposed to a low-flyer manager. Controls include: country, product category and year fixed effects, share of managers, number of blue-collar and white-collar workers.

Figure 1.20: Factory costs per ton and workers' past lateral moves

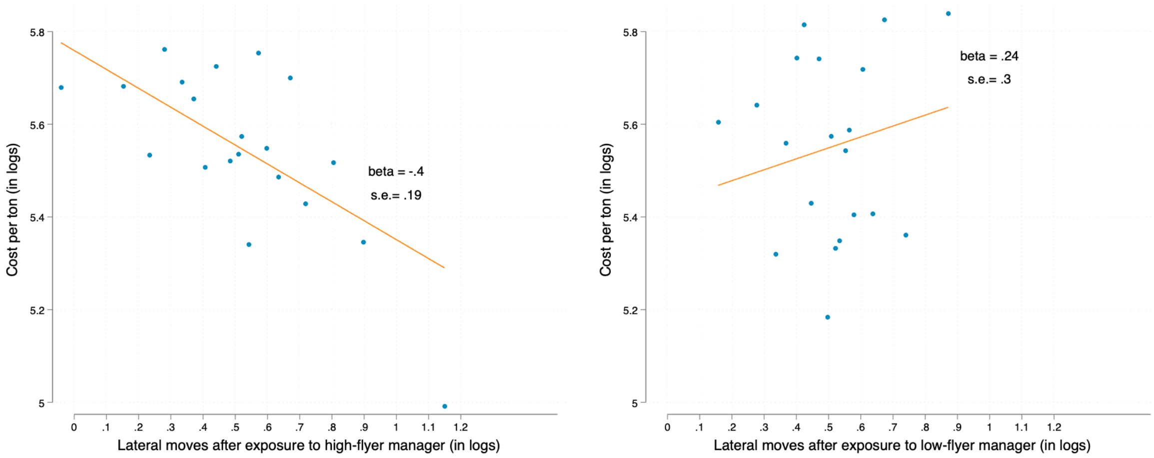

Notes. An observation is a factory-year. Because of changing reporting requirements, the costs per ton data could only be shared for the main product category (there are three product categories in total). Standard errors clustered by factory-year. The y-axis is operational costs per ton (EUR) in logs and the x-axis is the number of workers' lateral moves in logs: a) the left panel only considers the lateral moves that occur up to five years after being exposed to a high-flyer manager; b) the right panel only considers the lateral moves that occur up to five years after being exposed to a low-flyer manager. Controls include: country, and year fixed effects, share of managers, number of blue-collar and white-collar workers.

Figure 1.21: High-flyer manager and learning activities

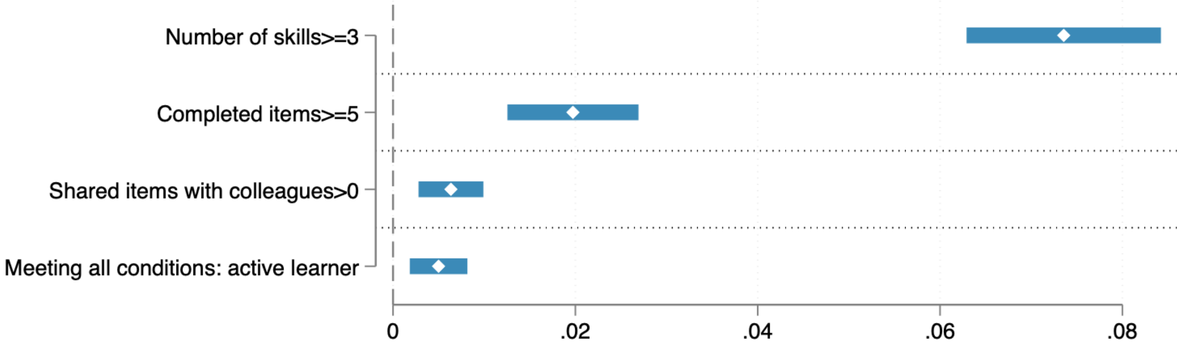

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Data taken from the career mobility platform at the firm introduced at the firm. Estimates obtained by running the model in equation 1.5. Outcome mean, low-flyer: Skills = 0.5; Items = 0.3 ; Shared = 0.04; Active learner= 0.04.

Figure 1.22: High-flyer manager and flexible projects

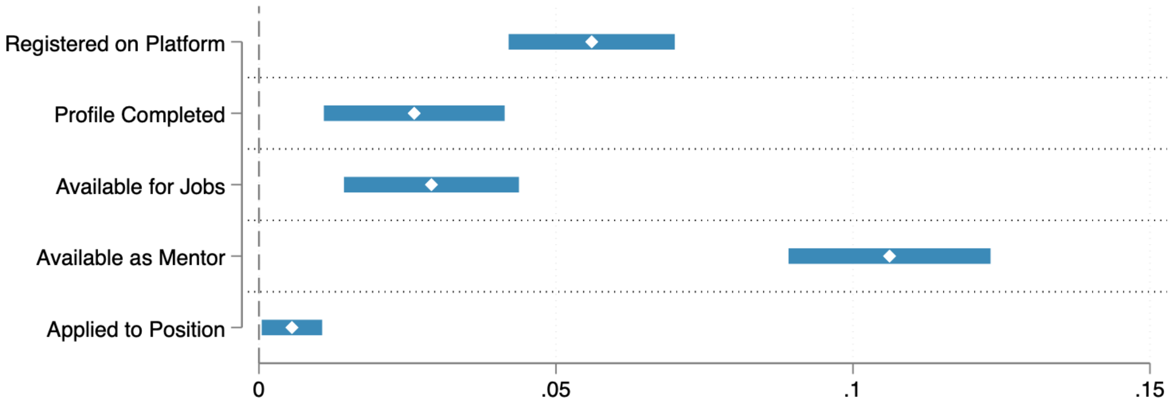

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Data taken from flexible project program at the firm introduced at the firm. Estimates obtained by running the model in equation 1.5. Outcome mean, low-flyer: Platform = 0.33; Profile = 0.47; Jobs = 0.25; Mentor = 0.12; Applied =0.02.

Figure 1.23: High-flyer manager and survey evidence

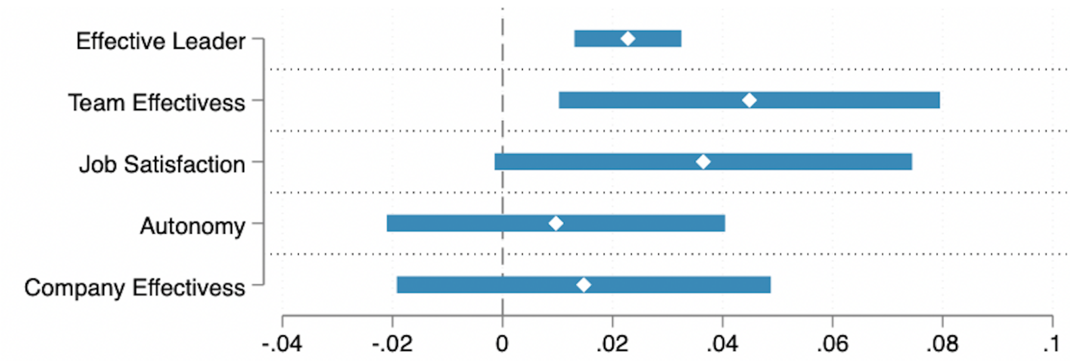

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Survey indices are the first principal components of various survey questions, grouped together by theme as detailed in Appendix Table 1.24. I use binary variables: probability of answering 5 out of 5-point Likert Scale. Estimates obtained by running the model in equation 1.5. Outcome mean, low-flyer (using the simple average for the indices): Effective Leader = 0.41 ; Team Effectiveness = 0.34; Job Satisfaction = 0.35; Autonomy = 0.32; Company effectiveness = 0.34. Appendix Figure 1.65 shows that the results are very similar when using simple averages for the indices instead of the first principal component.

Figure 1.24: Effects of gaining a high-flyer manager on moves within manager's network, probability

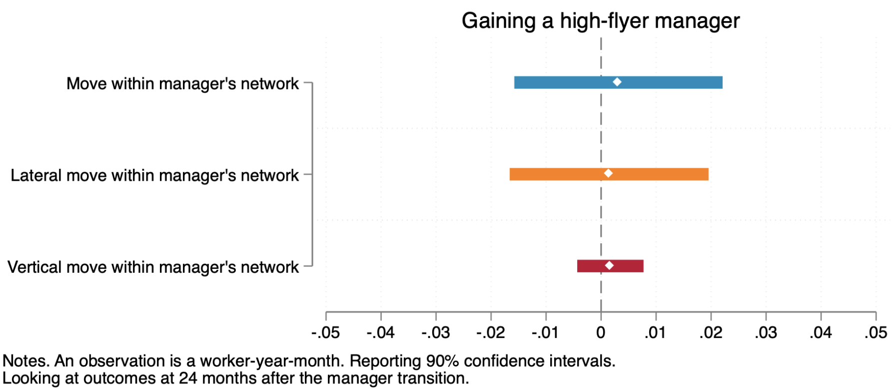

Notes. An observation is a worker-year-month. Looking up to 8 quarters since the manager transition (8 quarters is the average duration of a manager assignment to a team). 95% confidence intervals used and standard errors clustered by manager. Outcome mean for low-low transition: All moves= 19%; Lateral= 18%; Promotions=1%.

Figure 1.25: Changes in the organizational structure of team, jobs created and destroyed

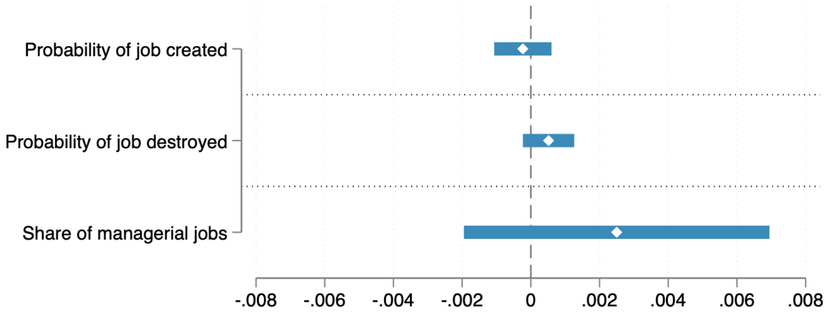

Notes. An observation is a team-year-month. 90% confidence intervals used and standard errors clustered by manager. Baseline Mean: New job= 0.015; Old job= 0.018; Share managerial jobs=0.17.

Figure 1.26: New hires: pay gap by high-flyer status of hiring manager

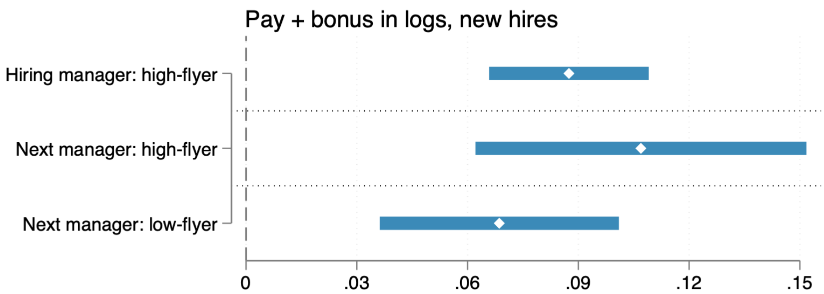

Notes. An observation is a worker-month. 90% confidence intervals used and standard errors clustered by manager. This figure shows the pay differences between workers by the manager type that hired them, four years after being hired, controlling for country and year-month fixed effect. The first row looks at all new hires; the second (third) row only considers new hires that move to a high (low)-flyer manager because of the managers' rotation policy.

Figure 1.27: Heterogeneous effects of gaining a high-flyer manager, worker and manager characteristics

(a) Manager tenure

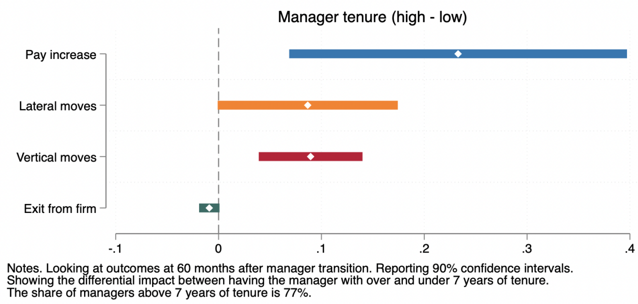

(b) Sharing gender with manager

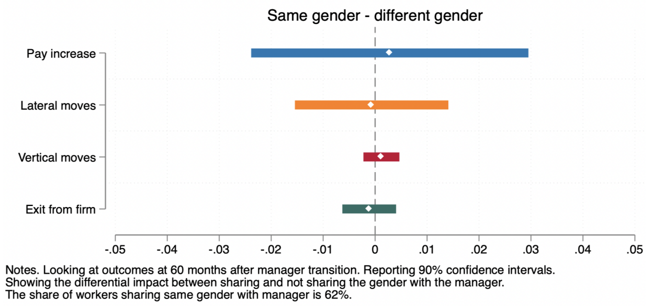

(c) Sharing office with manager

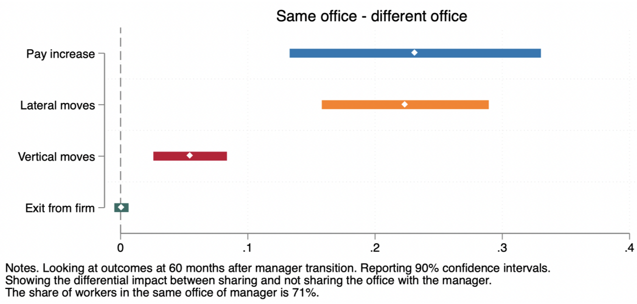

(d) Worker age

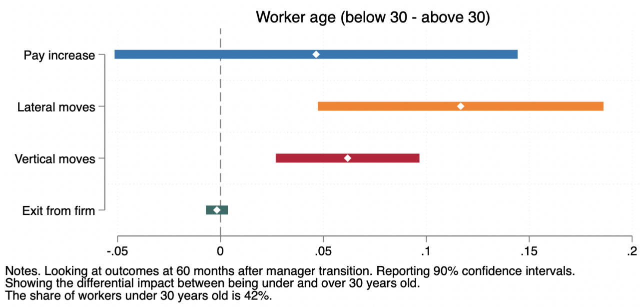

(e) Worker tenure

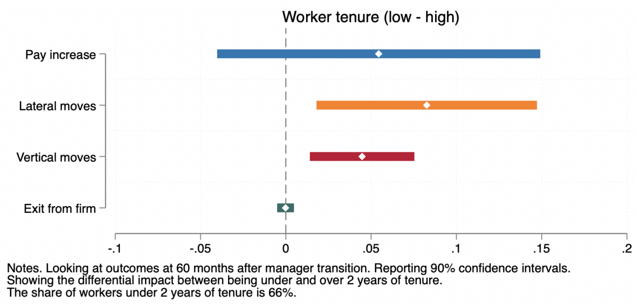

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Coefficients are estimated from a regression as in equation 1.7 and the figure reports the coefficient at the 20th quarter since the manager transition.

Figure 1.28: Heterogeneous effects of gaining a high-flyer manager, environment characteristics

(a) Office size

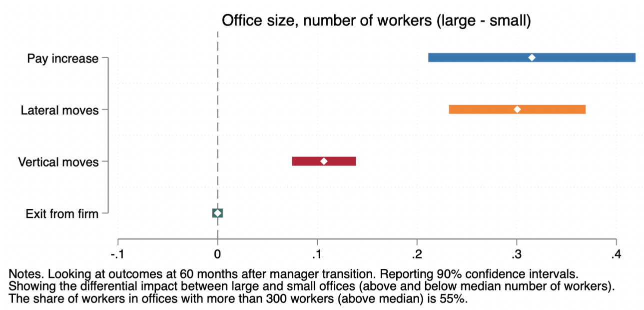

(b) Job diversity in the office

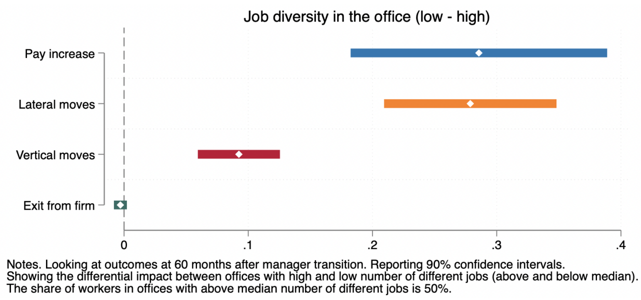

(c) Stringency of labor laws

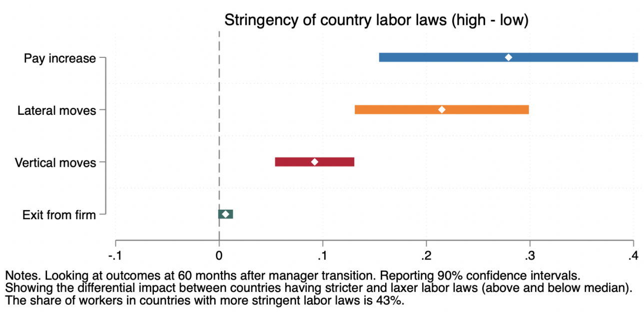

(d) Gender gap in countries with low FLFP

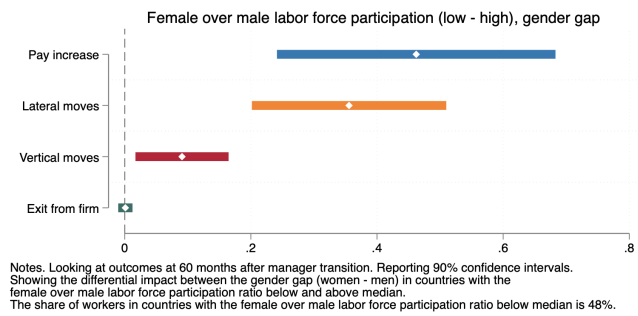

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Coefficients are estimated from a regression as in equation 1.7 and the figure reports the coefficient at the 20th quarter since the manager transition.

Figure 1.29: Heterogeneous effects of gaining a high-flyer manager, worker performance and moves

(a) Worker performance (p50)

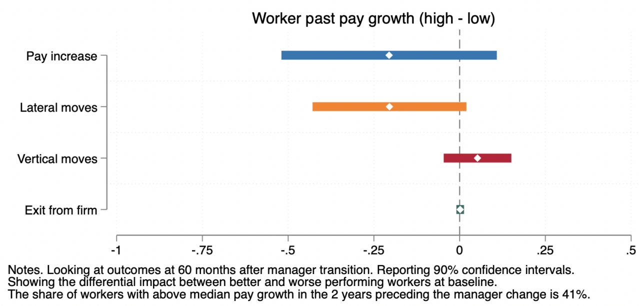

(b) Worker performance (p90 vs. p10)

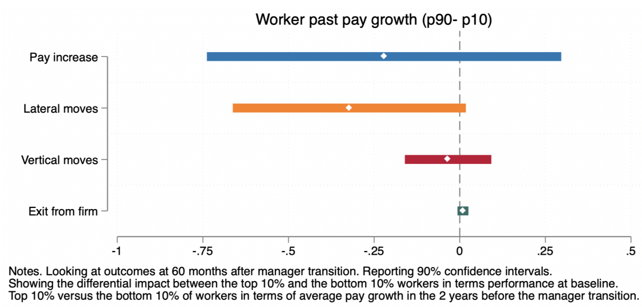

(c) Team performance (p50)

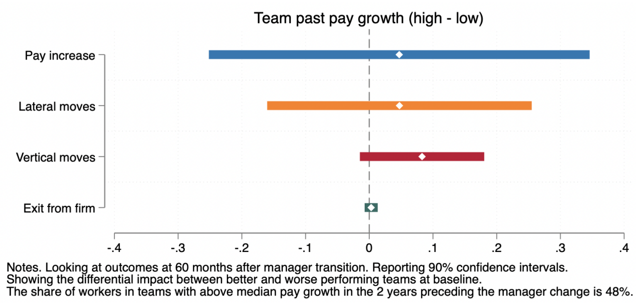

(d) Worker changes manager

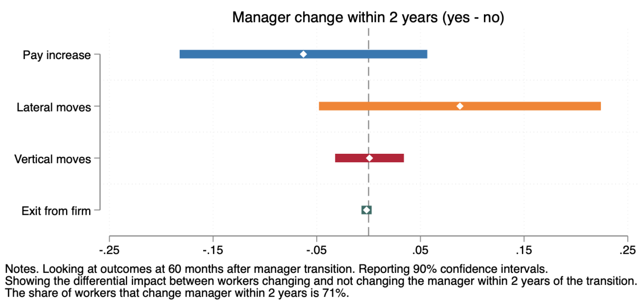

(e) Worker changes job

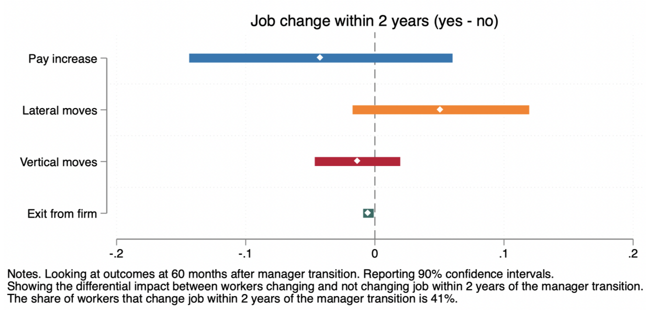

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Coefficients are estimated from a regression as in equation 1.7 and the figure reports the coefficient at the 20th quarter since the manager transition.

Figure 1.30: Effects of losing a high-flyer manager on lateral transfers, $(\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

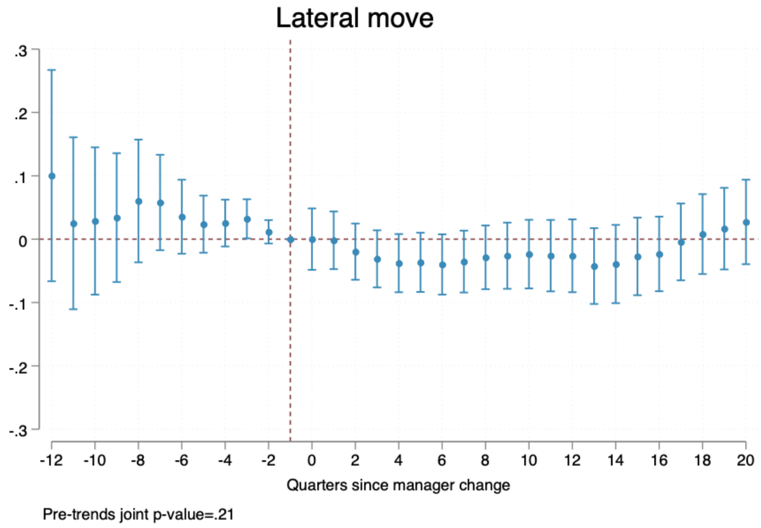

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The outcome variable is number of lateral transfers.

Figure 1.31: Effects of losing a high-flyer manager on cross-functional transfers, $(\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

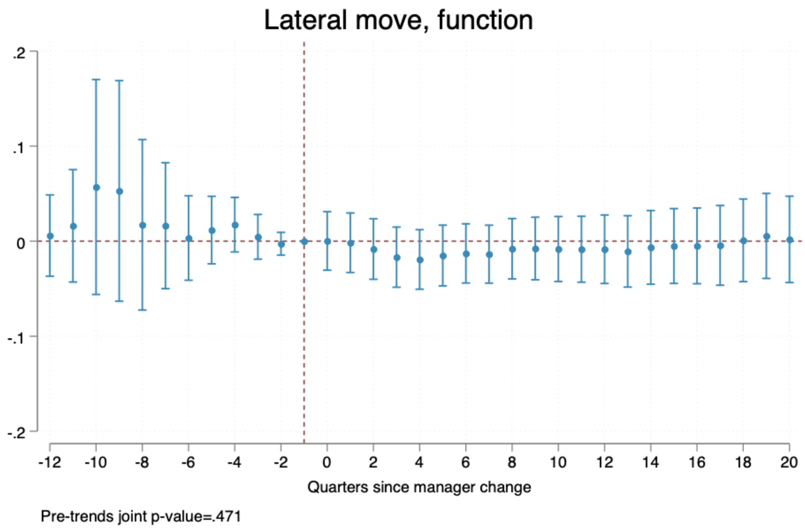

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. A cross-functional transfer is defined as a transfer across the 14 functions at the firm, e.g. from Finance to R&D.

Figure 1.32: Effects of losing a high-flyer manager on task distance in transfers (O*NET), $(\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

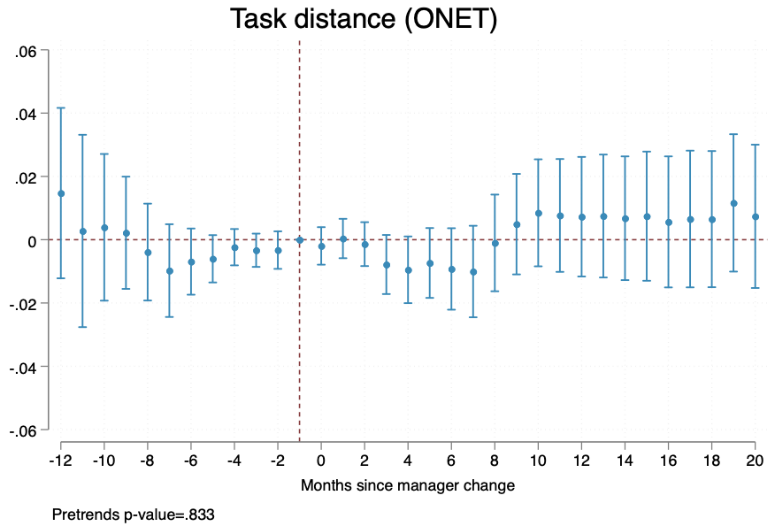

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. Task distance across jobs is constructed by matching the firm's job titles with O*NET data.

Figure 1.33: Effects of losing a high-flyer manager on work-level promotions, $(\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

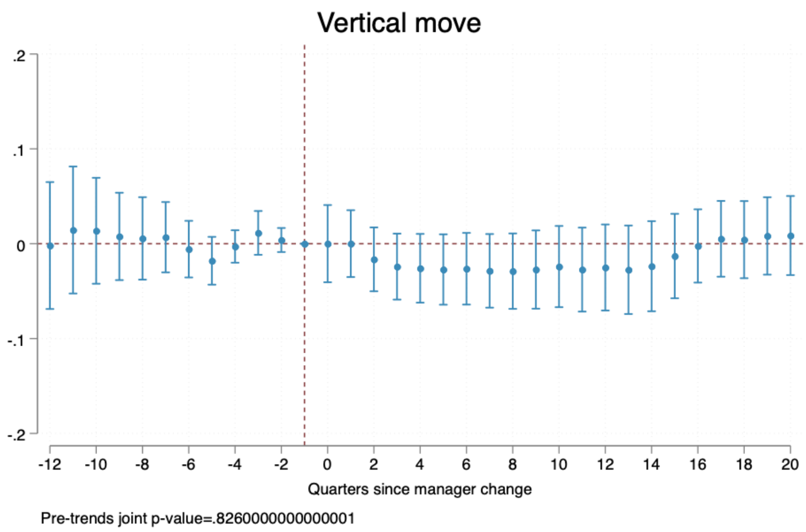

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The outcome variable is number of work-level promotions.

Figure 1.34: Effects of losing a high-flyer manager on exit from the firm, $(\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

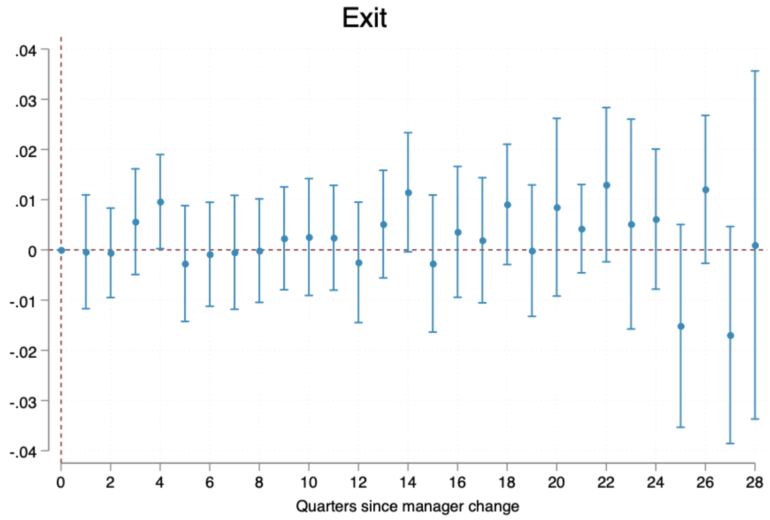

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The outcome variable is worker exit from the firm.

Figure 1.35: Effects of losing a high-flyer manager on quits and layoffs, $(\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

(a) Quits

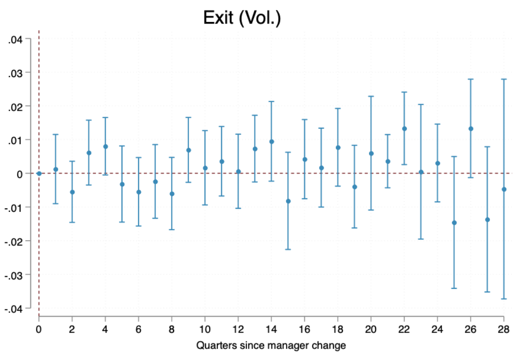

(b) Layoffs

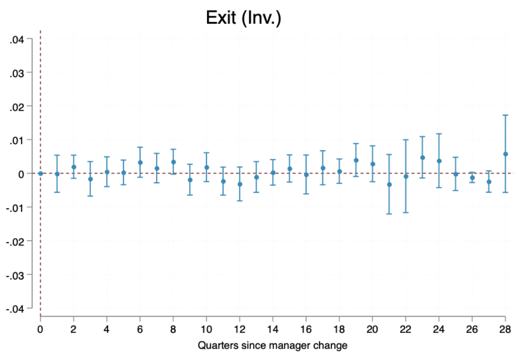

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The outcome variable is worker voluntary and involuntary exit from the firm.

Figure 1.36: Effects of losing a high-flyer manager on salary grade increases, $(\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

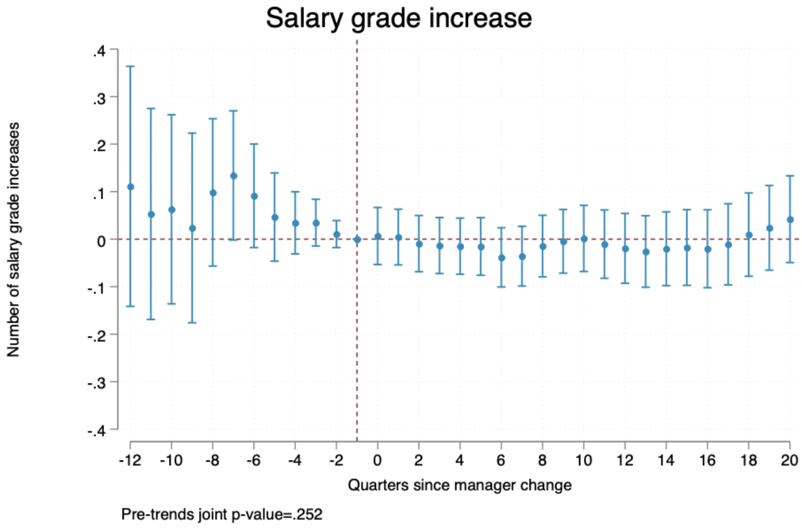

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The outcome variable is number of salary grade increases.

Figure 1.37: Effects of losing a high-flyer manager on total pay, $(\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

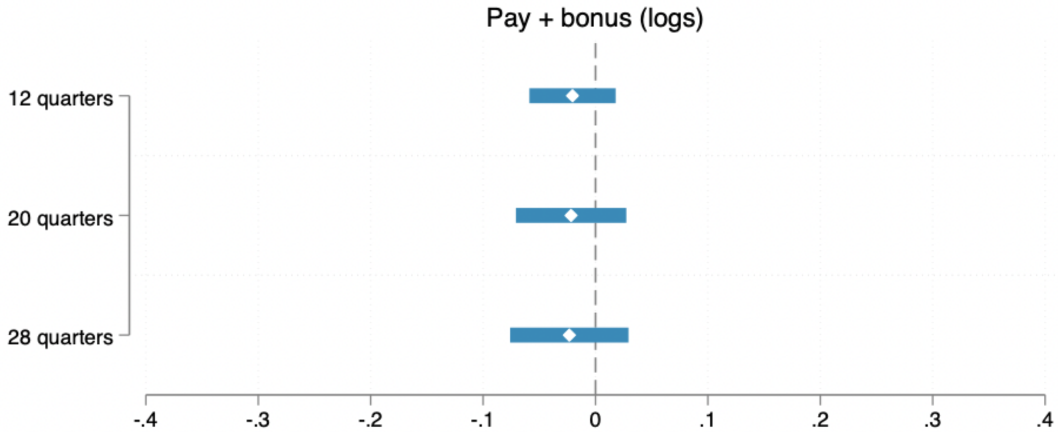

Notes. An observation is a worker-year-month. Aggregating the monthly coefficients to the quarterly level. Reporting the estimates at 12, 20 and 28 quarters after the manager transition. 90% confidence intervals used and standard errors clustered by manager.

Figure 1.38: Performance differential of workers promoted to managers, losing a high-flyer manager

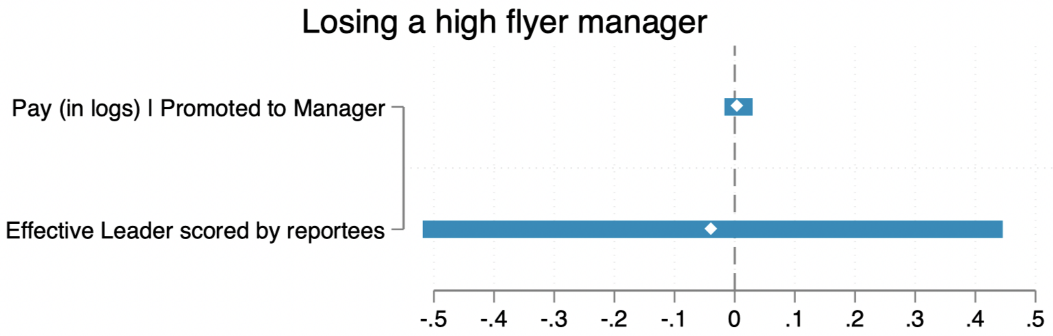

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Coefficients are estimated from separate regression as in equation 1.5. Sample restricted to the workers promoted to work-level 2.

Figure 1.39: Lateral transfer rates: test for asymmetries, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s}) - (\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

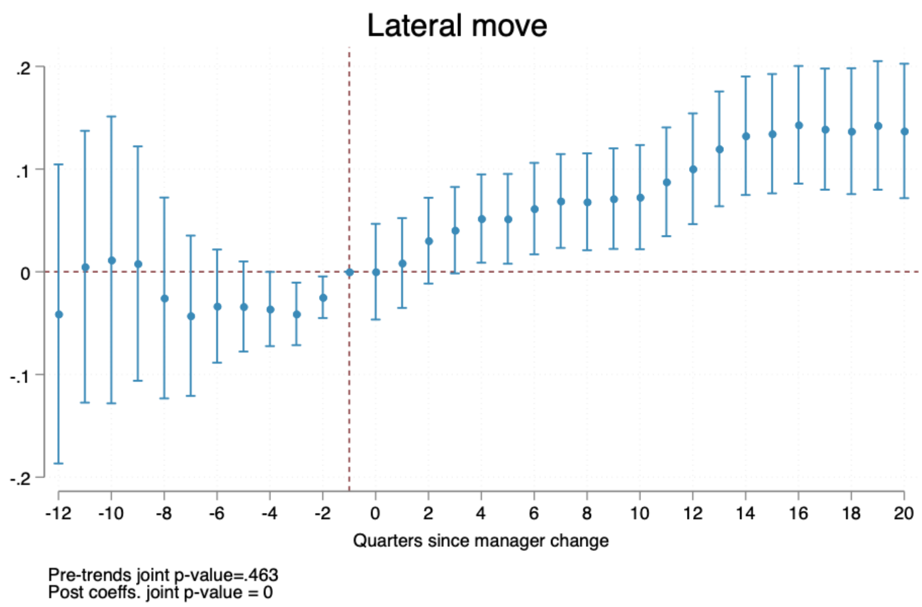

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager.

Figure 1.40: Cross-functional transfers: test for asymmetries, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s}) - (\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

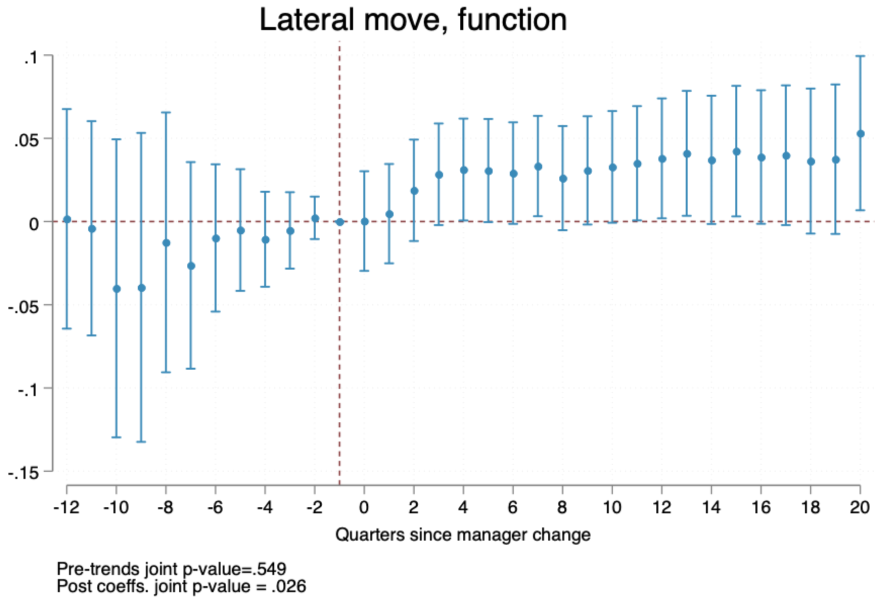

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager.

Figure 1.41: Vertical transfers: test for asymmetries, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s}) - (\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

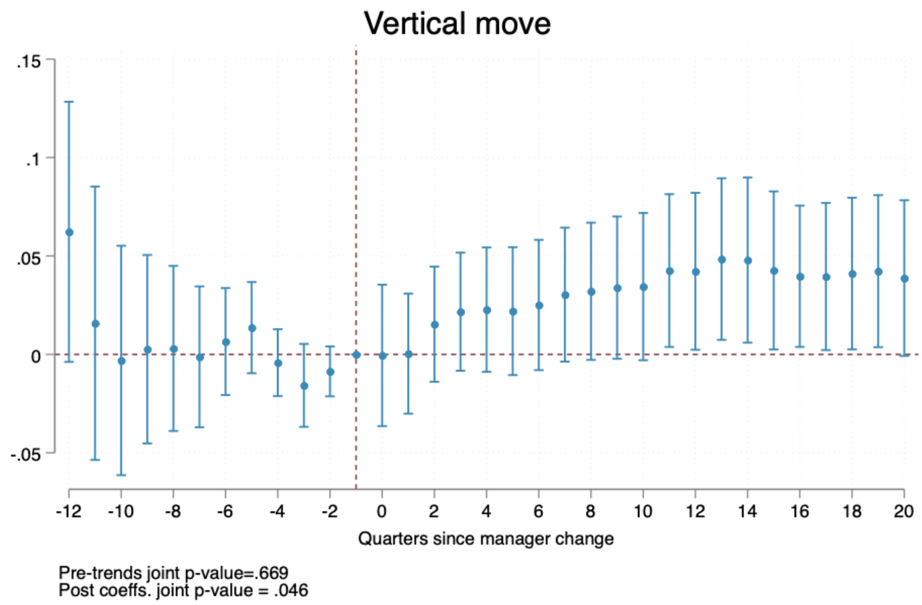

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager.

Figure 1.42: No differences in exit rates: test for asymmetries, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s}) - (\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

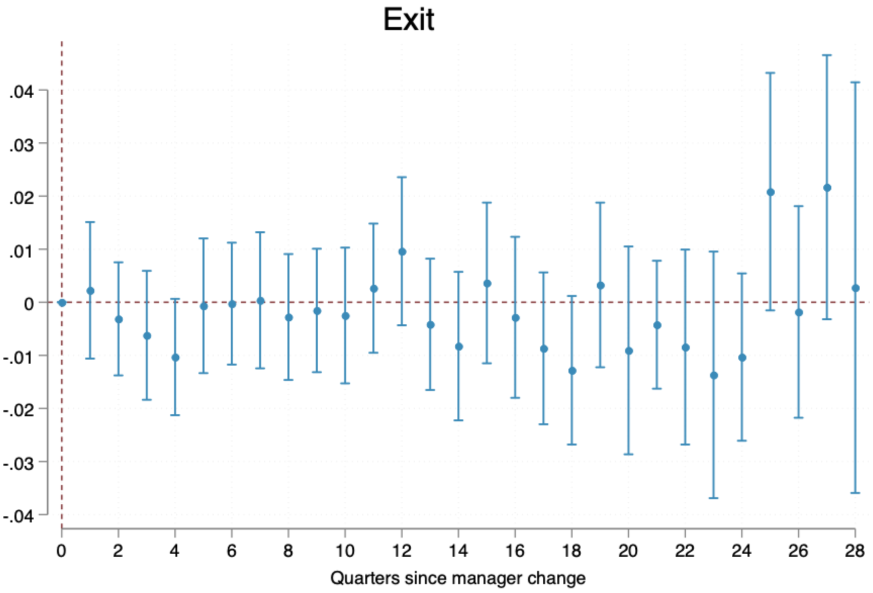

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager.

Figure 1.43: Pay (salary grade): test for asymmetries, $(\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s}) - (\hat{\beta}_{HtoL,s} - \hat{\beta}_{HtoH,s})$

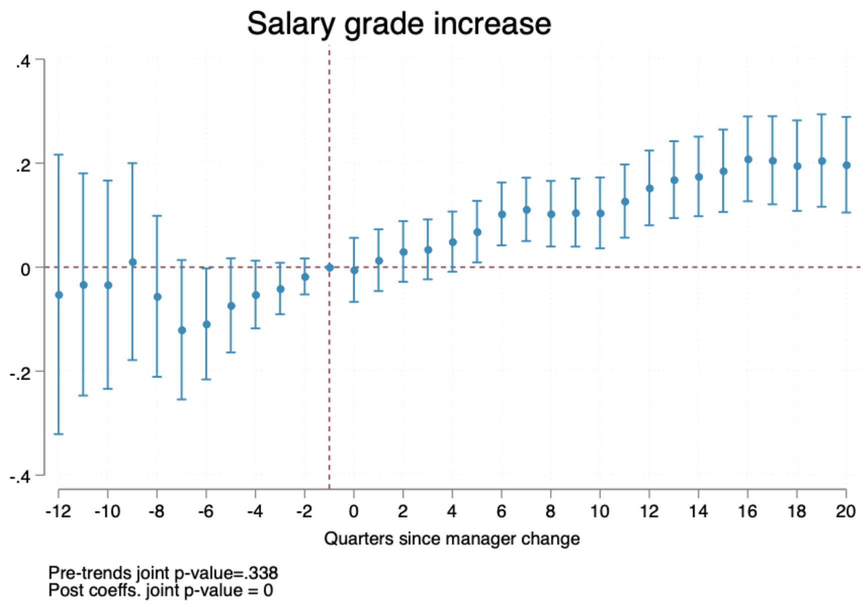

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager.

Figure 1.44: Effects of losing a high-flyer manager on moves within manager's network, probability

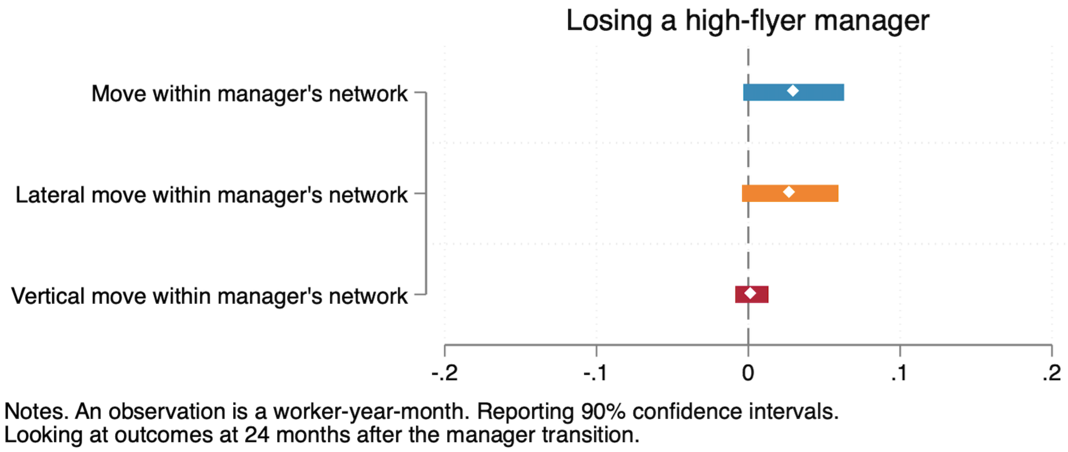

Notes. An observation is a worker-year-month. Looking up to 8 quarters since the manager transition (8 quarters is the average duration of a manager assignment to a team). 95% confidence intervals used and standard errors clustered by manager. Outcome mean for high-high transition: All moves= 17%; Lateral= 15%; Promotions=2%.

Figure 1.45: Losing a high-flyer manager: decomposing lateral moves

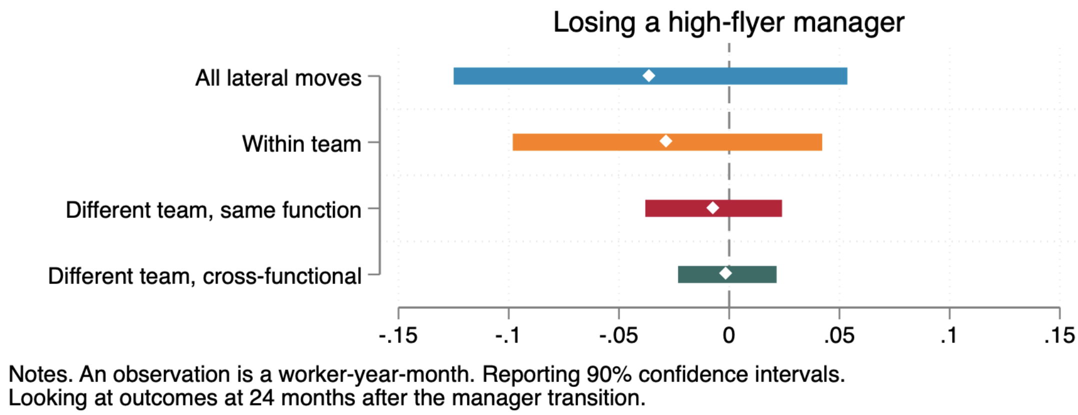

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Coefficients are estimated from a regression as in equation 1.2 and the figure reports the coefficient at the 8th quarter since the manager transition.

1.11 Tables

Table 1.1: Size of groups: workers, managers, jobs

Variable	**No. Unique Values**
Total white collar $\times$ months	8,618,267
Employee	205,432
Managers (work-level 2+)	30,511
Supervisors	42,744
Year- month	132
Standard Job	1,889
Sub-function $\times$ work-level	438
Offices	2,562
Countries	118
Country $\times$ Year	1,081
Office $\times$ Year	13,887
Employee $\times$ Job	423,644

Notes. An observation is a worker-month-year. The data contain personnel records for the entire white-collar employee base from January 2011 until December 2021.

Table 1.2: Descriptive Statistics

	Mean	SD	P1	P99	N
Panel A: Gender, age and education					
Female	0.43	0.5	0.0	1.0	8,616,744
Share in Cohort 18-29	0.26	0.4	0.0	1.0	8,618,267
Share in Cohort 30-39	0.39	0.5	0.0	1.0	8,618,267
Share in Cohort 40-49	0.23	0.4	0.0	1.0	8,618,267
Share in Cohort 50+	0.13	0.3	0.0	1.0	8,618,267
Econ, Business, and Admin	0.49	0.5	0.0	1.0	8,618,267
Sci, Engin, Math, and Stat	0.30	0.5	0.0	1.0	8,618,267
Social Sciences and Humanities	0.15	0.4	0.0	1.0	8,618,267
Other Educ	0.07	0.2	0.0	1.0	8,618,267
Panel B: Tenure, hierarchy and team size					
Tenure (years)	8.50	8.8	0.0	35.0	8,618,267
Share in Work-level 1	0.81	0.4	0.0	1.0	8,618,267
Share in Work-level 2	0.15	0.4	0.0	1.0	8,618,267
Share in Work-level 3+	0.04	0.2	0.0	1.0	8,618,267
No. of months per worker	41.95	36.5	1.0	111.0	205,432
No. of supervisors per worker	2.31	2.7	0.0	11.0	205,432
No. of workers per supervisor	5.07	7.9	1.0	33.0	42,744

Notes. An observation is a worker-month-year. The data contain personnel records for the entire white-collar employee base from January 2011 until December 2021. Cohort refers to the age group and work level denotes the hierarchical tier (from level 1 at the bottom to level 6).

Table 1.3: Descriptive Statistics, outcomes

	Mean	SD	P1	P99	N
Number of salary grade increases	0.60	1.0	0.0	4.0	224,117
Number of lateral job transfers	0.90	1.4	0.0	6.0	224,117
Number of promotions (work-level)	0.06	0.3	0.0	1.0	224,117
Monthly Exit	0.01	0.1	0.0	1.0	10083638
Pay + bonus (logs)	10.27	0.9	8.2	12.5	4,977,935
Bonus over Pay	0.20	116.2	0.0	0.6	4,977,935
Perf. appraisals	98.22	26.0	0.0	142.0	3,538,611
Productivity (sales in logs)	8.63	2.2	0.0	9.9	87,491

Notes. An observation is a worker-month-year. The data contain personnel records for the entire white-collar employee base from January 2011 until December 2021. Salary information is only available since January 2015 and the data on performance appraisals start in January 2017.

Table 1.4: High-flyer manager: performance after high-flyer status is determined

	(1)	(2)	(3)
Variable	Not High Flyer	High Flyer	Difference
Salary growth	0.006	0.011	0.005***
	(0.029)	(0.034)	(0.000)
Promotion work-level 3	0.049	0.135	0.086***
	(0.204)	(0.331)	(0.000)
Perf. appraisal (1-150)	100.932	104.042	3.110***
	(18.236)	(15.963)	(0.000)
Effective leader (survey)	4.025	4.119	0.095***
	(0.691)	(0.681)	(0.000)
Observations	12,385	6,875	19,298

Notes. Showing mean and standard deviations (in parentheses) and p-values for the difference in means. The difference in means is computed using standard errors clustered by manager.

Table 1.5: High-flyer manager: demographics

Variable	(1) Not High Flyer	(2) High Flyer	(3) Difference
Female	0.498	0.565	0.067***
	(0.500)	(0.496)	(0.000)
MBA	0.001	0.001	-0.000
	(0.026)	(0.026)	(0.981)
Econ, Business, and Admin	0.504	0.553	0.050***
	(0.500)	(0.497)	(0.008)
Sci, Engin, Math, and Stat	0.274	0.262	-0.012
	(0.446)	(0.440)	(0.470)
Social Sciences and Humanities	0.163	0.165	0.002
	(0.369)	(0.371)	(0.881)
Other Educ	0.065	0.029	-0.036***
	(0.247)	(0.168)	(0.000)
Mid-career recruit	0.314	0.133	-0.181***
	(0.464)	(0.339)	(0.000)
Hired through graduate programme	0.031	0.356	0.325***
	(0.174)	(0.479)	(0.000)
Observations	13,979	5,636	19,615

Notes. Showing mean and standard deviations (in parentheses) and p-values for the difference in means. The difference in means is computed using standard errors clustered by manager.

Table 1.6: Endogenous mobility check: team performance

	(1)	(2)	(3)	(4)	(5)
	Prom. (salary)	Prom. (work level)	Pay Growth	Perf. Appraisals (CV)	Lateral job change
High-flyer manager	0.000422	0.000216	0.00203	-0.000986	0.00192*
	(0.000684)	(0.000164)	(0.00155)	(0.0147)	(0.00103)
Controls	Yes	Yes	Yes	Yes	Yes
Team FE	No	No	No	No	No
Mean	0.0147	0.000763	-0.000439	0.171	0.0214
N	112414	112414	23278	8604	112414
R-squared	0.00767	0.00232	0.00858	0.176	0.0116

Notes. An observation is a team-month. Sample restricted to observations between 6 and 36 months before the manager switch. Controls include: function, country and year FE, and team size. Standard errors clustered at the manager level.

Table 1.7: Endogenous mobility check: team diversity

	(1)	(2)	(3)	(4)
	Diversity, gender	Diversity, age	Diversity, office	Diversity, nationality
High-flyer manager	0.00436	-0.000149	-0.00282	-0.00499
	(0.00524)	(0.00505)	(0.00851)	(0.00427)
Controls	Yes	Yes	Yes	Yes
Team FE	No	No	No	No
Mean	0.291	0.492	0.195	0.0772
N	112414	112414	112414	112414
R-squared	0.106	0.146	0.154	0.174

Notes. An observation is a team-month. Sample restricted to observations between 6 and 36 months before the manager switch. Controls include: function, country and year FE, and team size. Standard errors clustered at the manager level.

Table 1.8: Endogenous mobility check: team homophily with manager

	(1) Same Gender	(2) Same Age Band	(3) Same Office	(4) Same Nationality
High-flyer manager	-0.0148	0.00516	0.00482	0.00468
	(0.0110)	(0.00933)	(0.0111)	(0.00667)
Controls	Yes	Yes	Yes	Yes
Team FE	No	No	No	No
Mean	0.632	0.305	0.788	0.915
N	112411	112414	112414	112113
R-squared	0.0654	0.0299	0.150	0.165

Notes. An observation is a team-month. Sample restricted to observations between 6 and 36 months before the manager switch. Controls include: function, country and year FE, and team size. Standard errors clustered at the manager level.

Table 1.9: Endogenous mobility check: team performance (asymmetric)

	(1) Prom. (salary)	(2) Prom. (work level)	(3) Pay Growth	(4) Perf. Appraisals (CV)	(5) Lateral job change
Low to High	0.000447	0.000292	0.00193	-0.00426	0.00229**
	(0.000762)	(0.000189)	(0.00176)	(0.0168)	(0.00116)
High to High	0.00306**	0.0000733	0.00424	-0.00140	0.00155
	(0.00125)	(0.000269)	(0.00266)	(0.0234)	(0.00170)
High to Low	0.00299***	0.000315	0.00115	-0.0116	0.00200*
	(0.000850)	(0.000199)	(0.00206)	(0.0162)	(0.00121)
Controls	Yes	Yes	Yes	Yes	Yes
Team FE	No	No	No	No	No
Mean	0.0147	0.000763	-0.000439	0.171	0.0214
N	112414	112414	23278	8604	112414
R-squared	0.00781	0.00235	0.00860	0.176	0.0116
HtoL - HtoH	-0.0000626	0.000242	-0.00309	-0.0102	0.000447
p-value:	0.964	0.444	0.302	0.682	0.820

Notes. An observation is a team-month. Sample restricted to observations between 6 and 36 months before the manager switch. Controls include: function, country and year FE, and team size. Standard errors clustered at the manager level.

Table 1.10: Endogenous mobility check: team diversity (asymmetric)

	(1)	(2)	(3)	(4)
	Diversity, gender	Diversity, age	Diversity, office	Diversity, nationality
Low to High	0.00676	-0.00215	-0.00648	-0.00440
	(0.00581)	(0.00558)	(0.00929)	(0.00469)
High to High	-0.00689	-0.0141	0.00231	-0.00282
	(0.00909)	(0.00897)	(0.0163)	(0.00711)
High to Low	0.00278	-0.0214***	-0.000567	0.00236
	(0.00678)	(0.00649)	(0.0106)	(0.00578)
Controls	Yes	Yes	Yes	Yes
Team FE	No	No	No	No
Mean	0.291	0.492	0.195	0.0772
N	112414	112414	112414	112414
R-squared	0.106	0.147	0.154	0.174
HtoL - HtoH	0.00967	-0.00733	-0.00288	0.00518
p-value:	0.349	0.477	0.873	0.543

Notes. An observation is a team-month. Sample restricted to observations between 6 and 36 months before the manager switch. Controls include: function, country and year FE, and team size. Standard errors clustered at the manager level.

Table 1.11: Endogenous mobility check: team homophily with manager (asymmetric)

	(1) Same Gender	(2) Same Age Band	(3) Same Office	(4) Same Nationality
Low to High	-0.00508	-0.00248	0.00945	0.00295
	(0.0123)	(0.0101)	(0.0122)	(0.00768)
High to High	-0.0686***	0.0418**	-0.0119	0.00576
	(0.0189)	(0.0178)	(0.0212)	(0.0100)
High to Low	-0.0431***	0.0439***	-0.00678	0.00603
	(0.0142)	(0.0120)	(0.0141)	(0.00822)
Controls	Yes	Yes	Yes	Yes
Team FE	No	No	No	No
Mean	0.632	0.305	0.788	0.915
N	112411	112414	112414	112113
R-squared	0.0677	0.0321	0.150	0.165
HtoL - HtoH	0.0256	0.00206	0.00509	0.000269
p-value:	0.243	0.919	0.829	0.982

Notes. An observation is a team-month. Sample restricted to observations between 6 and 36 months before the manager switch. Controls include: function, country and year FE, and team size. Standard errors clustered at the manager level.

Table 1.12: Characteristics of employees by type of manager transition

Variable	(1) No event	(2) Had event	(3) Low to Low	(4) Low to High	(5) High to Low	(6) High to High
Female	0.437	0.461	0.452	0.475	0.489	0.490
	(0.496)	(0.498)	(0.498)	(0.499)	(0.500)	(0.500)
Age	37.257	33.679	34.546	32.242	31.223	30.801
	(9.020)	(8.518)	(8.780)	(7.747)	(7.194)	(7.107)
MBA	0.000	0.000	0.000	0.000	0.000	0.000
	(0.017)	(0.015)	(0.018)	(0.000)	(0.000)	(0.000)
Econ, Business, and Admin	0.480	0.469	0.437	0.516	0.551	0.577
	(0.500)	(0.499)	(0.496)	(0.500)	(0.498)	(0.495)
Sci, Engin, Math, and Stat	0.307	0.303	0.335	0.256	0.236	0.187
	(0.461)	(0.460)	(0.472)	(0.437)	(0.426)	(0.391)
Social Sciences and Humanities	0.146	0.144	0.137	0.146	0.155	0.198
	(0.353)	(0.352)	(0.344)	(0.354)	(0.363)	(0.399)
Other Educ	0.073	0.087	0.094	0.087	0.064	0.049
	(0.260)	(0.281)	(0.291)	(0.282)	(0.246)	(0.217)
Tenure (years)	8.535	5.172	5.998	3.711	2.766	2.750
	(8.816)	(7.944)	(8.391)	(6.714)	(5.986)	(5.928)
No. of job moves	1.067	0.207	0.221	0.199	0.144	0.166
	(1.450)	(0.544)	(0.565)	(0.531)	(0.424)	(0.477)
No. of salary increases	0.610	0.190	0.214	0.160	0.110	0.105
	(0.938)	(0.467)	(0.496)	(0.425)	(0.339)	(0.337)
Pay growth (6 months)	0.017	0.009	0.006	0.005	0.015	0.031
	(0.260)	(0.228)	(0.231)	(0.187)	(0.270)	(0.208)
Perf. appraisal (1-150)	98.249	97.014	96.986	97.027	96.847	97.440
	(26.004)	(25.620)	(25.359)	(26.641)	(26.155)	(24.602)
Unique num. workers, work level 1	205432	80687	60614	9466	7195	3412

Notes. This table shows summary statics across event types, and between the groups who do and do not experience events. Outgoing managers are defined as the manager of unit in the month before a transition event; incoming managers are those who are assigned to a unit in the month of the event. For event columns, I show the average of employees in the month they experience the event; for those who never experience an event I show the average of all such individuals across their tenure at the firm.

1.12 Appendix Figures

Figure 1.46: Data sources and time period

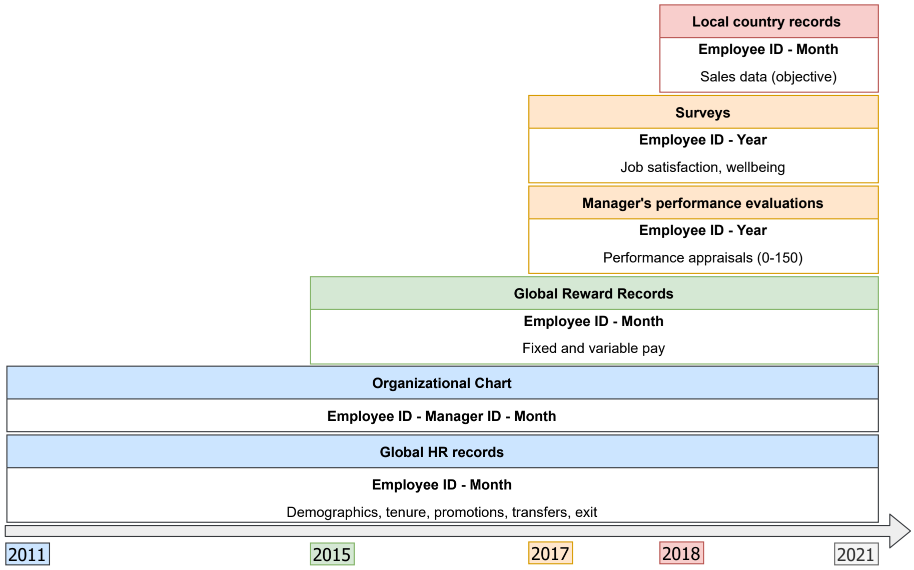

Notes. This figure shows the data sources collated from the multinational's records.

Figure 1.47: Hierarchy

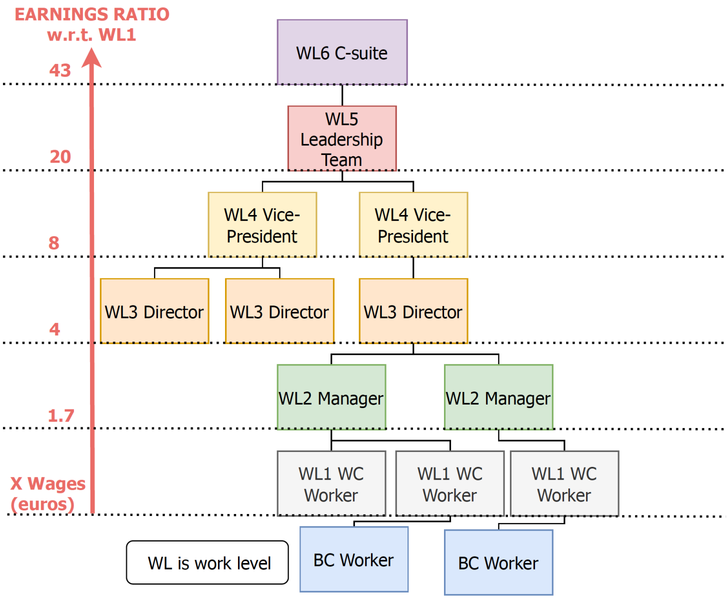

Notes. This figure shows the vertical job differentiation at the company.

Figure 1.48: Profile screen in the career mobility platform

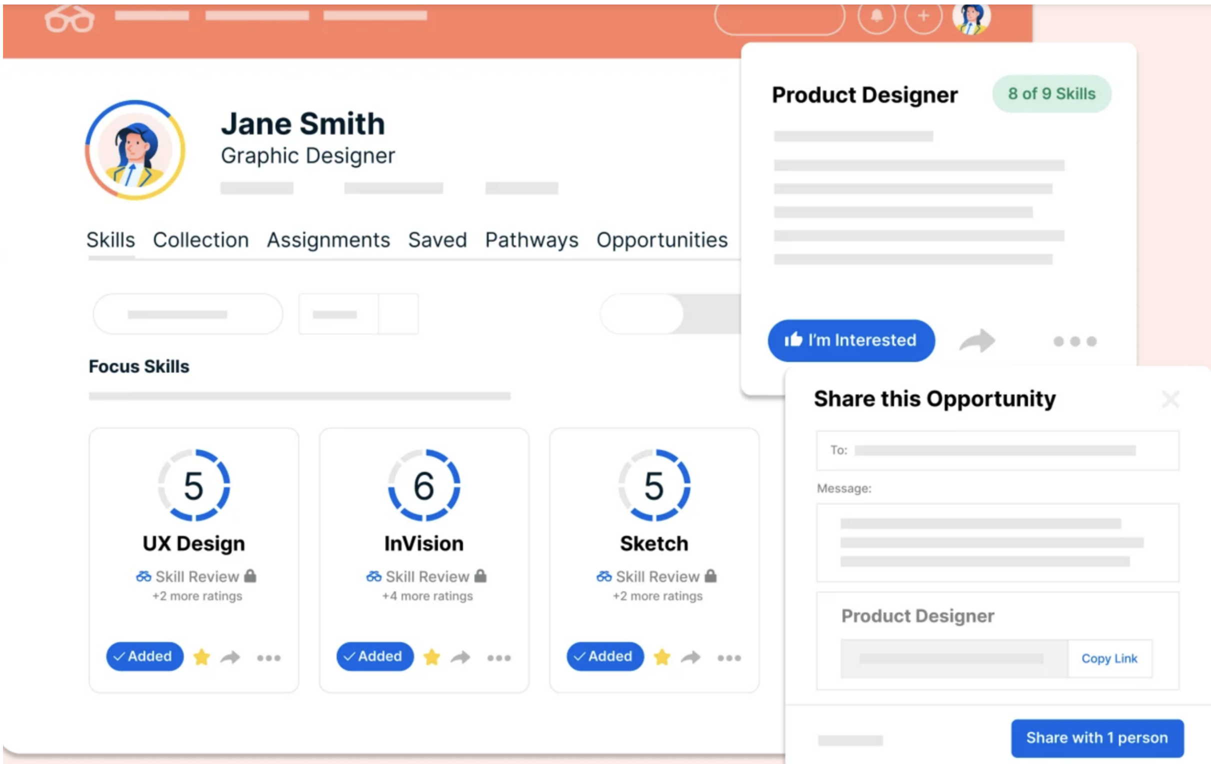

Notes. This figure portrays a mock profile screen of an employee in the career mobility platform. It is a talent tool provided by a third-party that combines learning and development, skill analytics and career mobility.

Figure 1.49: Guidelines for line managers

Notes. This figure is an excerpt from the guidelines set by HR for managers regarding the contents of the check-ins managers should be doing with their teams on a weekly basis.

Figure 1.50: Word cloud of skills sought after by employees

Notes. This word cloud pictures the focus skills listed in the career mobility platform by employees in the firm. This data is only available for a random subset of around 3000 employees.

Figure 1.51: Standard deviation of pay (within job-office-month)

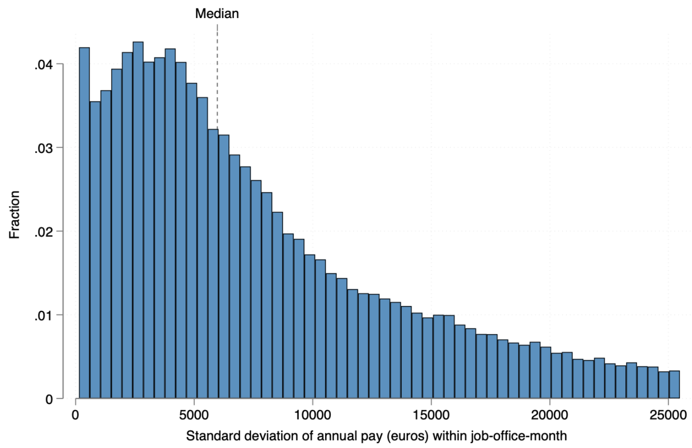

Notes. This graph shows the distribution of the standard deviation in overall pay (fixed pay plus variable pay) within a given job title in an office and year-month.

Figure 1.52: Average job transfer rates by sub-function

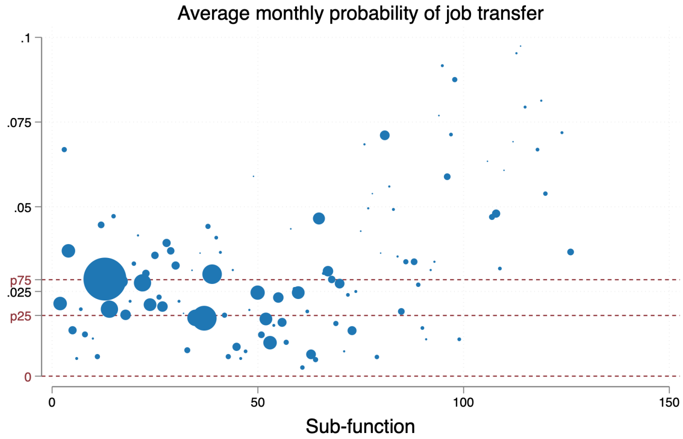

Notes. This figure shows the average monthly probability of a lateral move by sub-function. The size of the circles is proportional to the size of the sub-function.

Figure 1.53: Average salary increase rates by sub-function

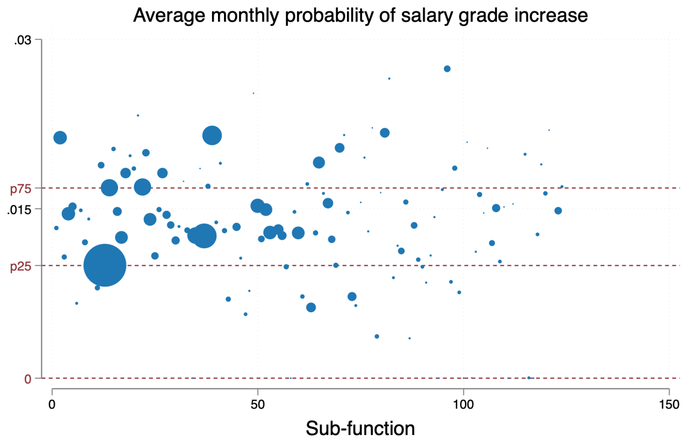

Notes. This figure shows the average monthly probability of a salary grade increase by sub-function. The size of the circles is proportional to the size of the sub-function.

Figure 1.54: Tenure profiles over the years, by work-level

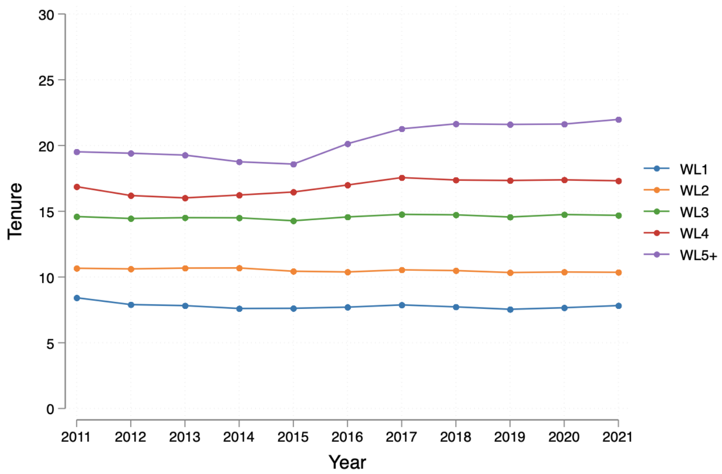

Notes. This figure shows the average tenure across work-levels over the years.

Figure 1.55: Age profiles over the years, by work-level

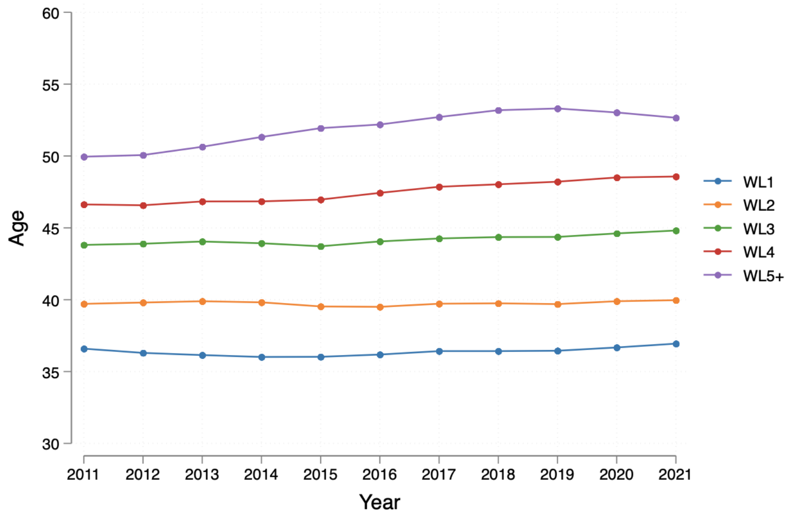

Notes. This figure shows the average age across work-levels over the years. Since the data comes aggregated into 10-year age groups, I take the age in the middle to create a continuous variable for age.

Figure 1.56: Work-level shares over the years, by work-level

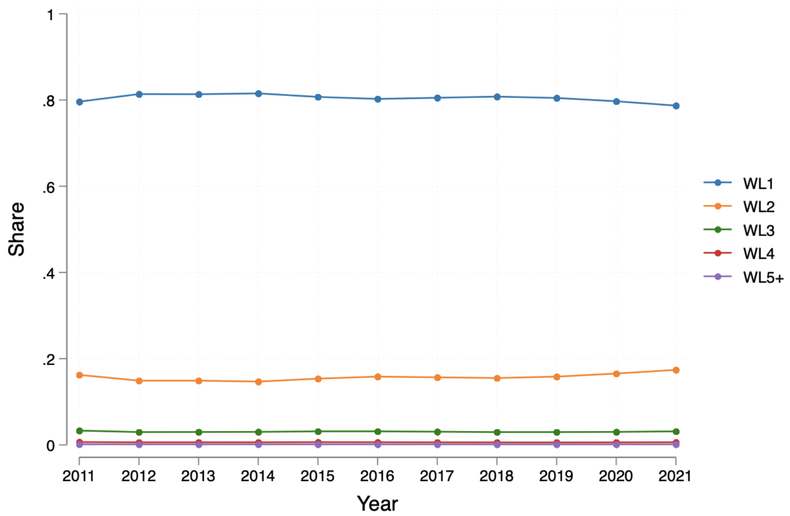

Notes. This figure shows the share of workers across work-levels over the years.

Figure 1.57: Empirical Bayes shrunk manager fixed effect on worker future pay (in logs)

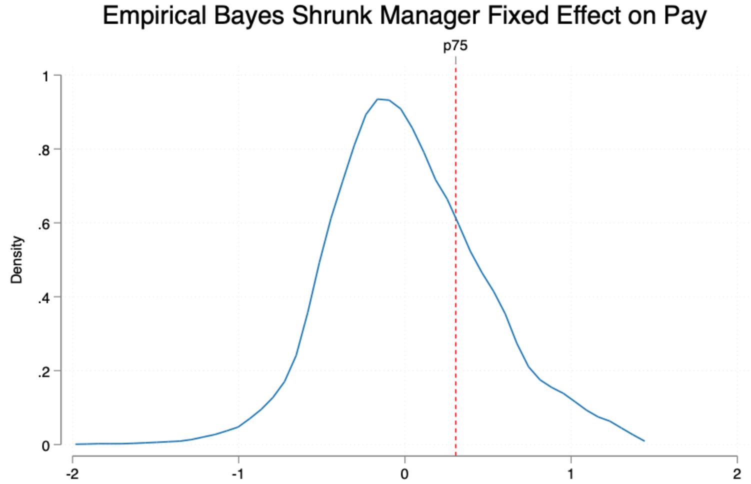

Notes. I take worker pay in logs with a five year gap with respect to the manager exposure. I perform an Empirical Bayes shrinkage procedure to account for the upward bias in the variance due to sampling noise (Morris, 1983).

Figure 1.58: CDF of the duration of managers' previous job before a new transition

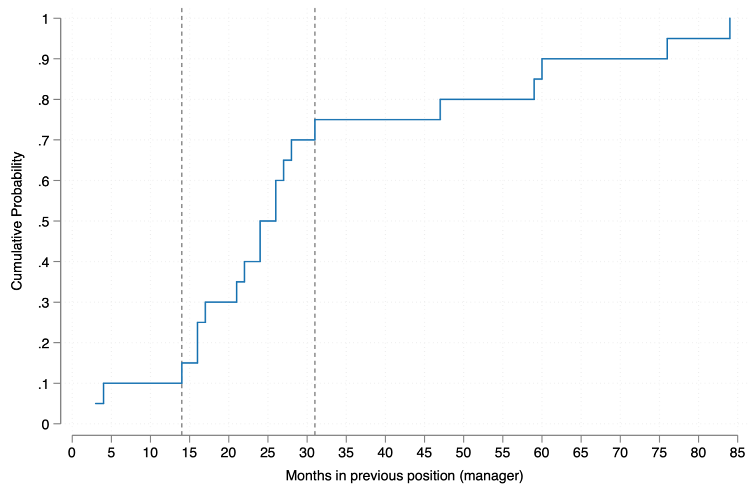

Notes. Cumulative distribution function of the number of months in the job before a manager makes the next team transition.

Figure 1.59: Relationship between productivity in standard deviations and future promotion

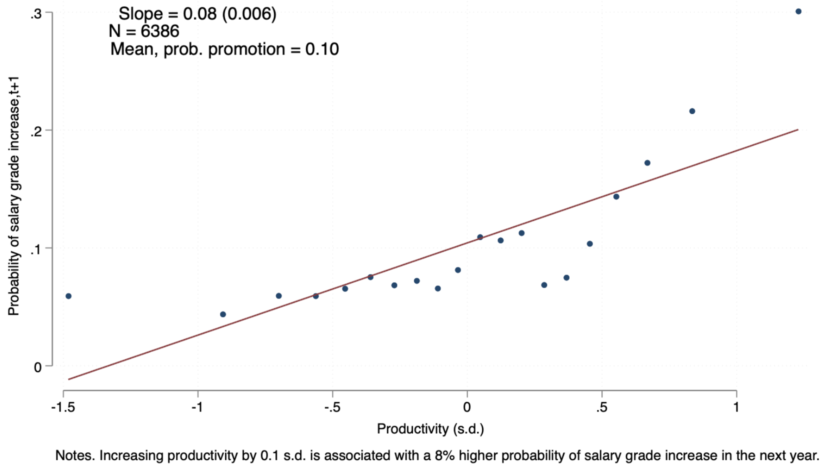

Notes. The above presents binned scatter plots of objective sales bonus against the probability of a salary grade increase the next year. Sample of Indian sales workers, 2018-2021.

Figure 1.60: Relationship between productivity in 1000s Indian rupees and future promotion

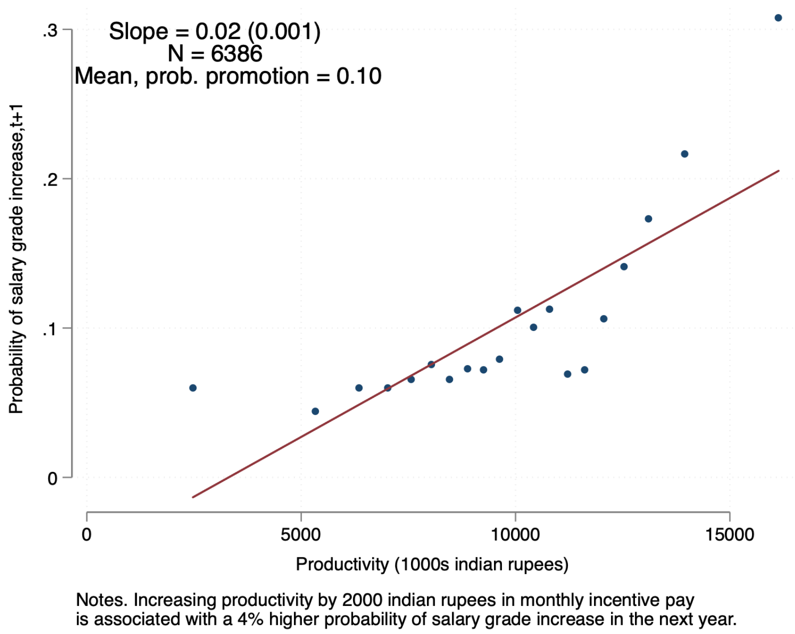

Notes. The above presents binned scatter plots of objective sales bonus against the probability of a salary grade increase the next year. Sample of Indian sales workers, 2018-2021.

Figure 1.61: Relationship between pay and salary grade increases

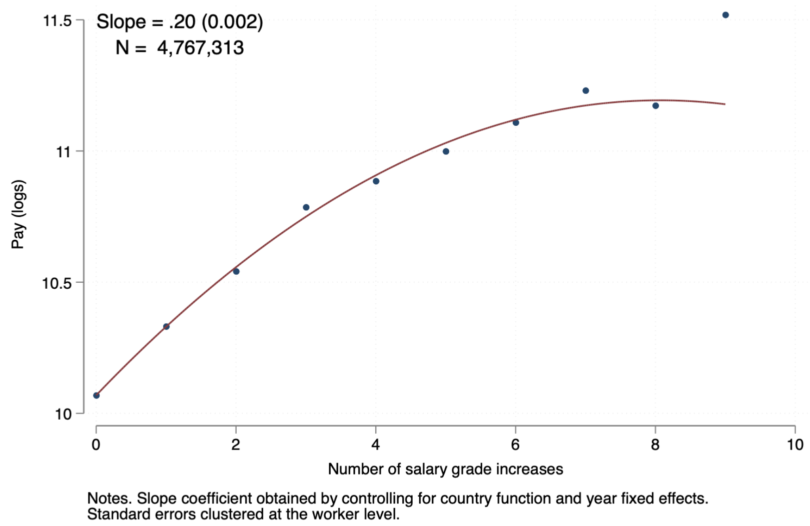

Notes. The above presents a binned scatter plot of pay and number of salary grade increases.

Figure 1.62: Performance differential of workers promoted to managers

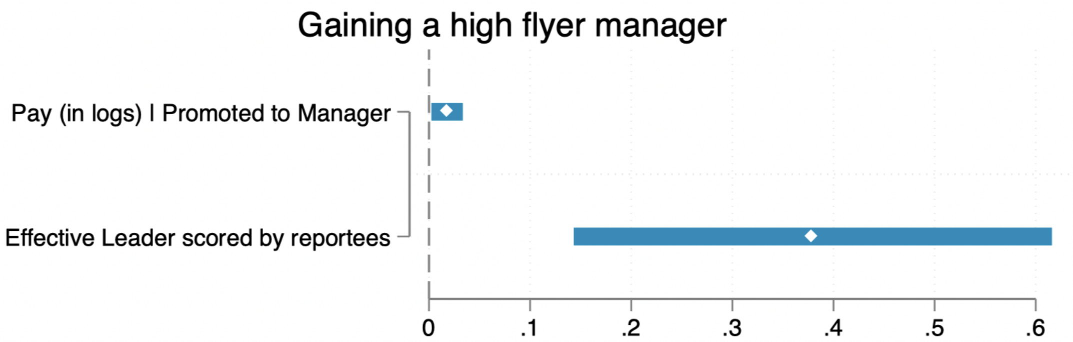

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Coefficients are estimated from separate regression as in equation 1.5. Sample restricted to the workers promoted to work-level 2.

Figure 1.63: Voluntary exit survey: the effect of high-flyer manager on reason for changing job

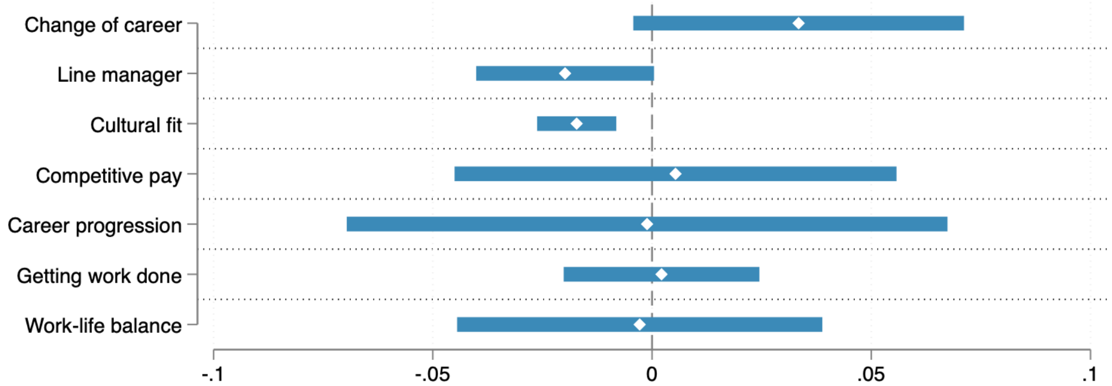

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Data comes from the voluntary exit survey that workers who quit the organization are invited to participate in. Showing the $\hat{\alpha}_1$ coefficient obtained from running this model: $y_{it} = \alpha_0 + \alpha_1 \text{High Flyer}_{it} + \boldsymbol{X_{it}}'\boldsymbol{\beta} + \eta_{it}$.

Figure 1.64: Effects of gaining a high-flyer manager on fixed pay ($\hat{\beta}_{LtoH,s} - \hat{\beta}_{LtoL,s}$)

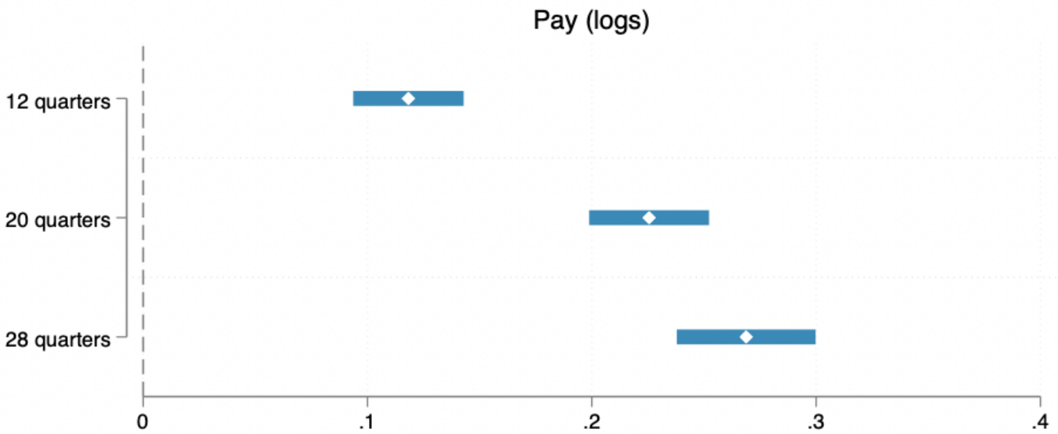

Notes. An observation is a worker-year-month. Aggregating the monthly coefficients to the quarterly level. Reporting the estimates at 12, 20 and 28 quarters after the manager transition. 90% confidence intervals used and standard errors clustered by manager.

Figure 1.65: High-flyer manager and survey evidence: using averages for the indices

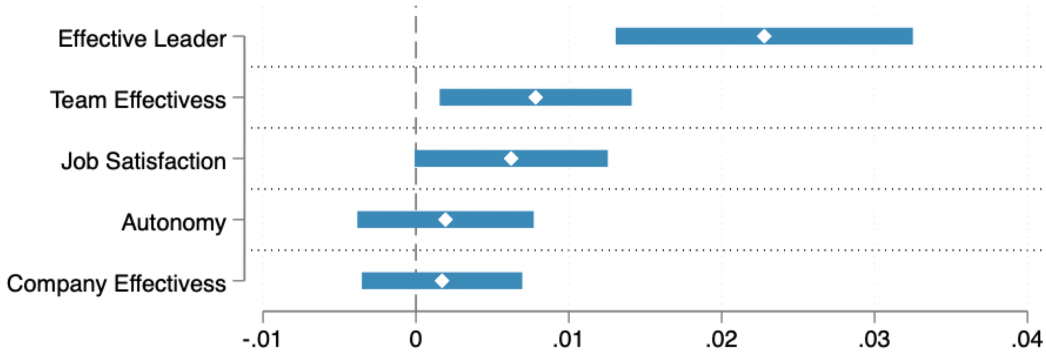

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Survey indices are the first principal components of various survey questions, grouped together by theme as detailed in Appendix Table 1.24. I use binary variables: probability of answering 5 out of 5-point Likert Scale. Estimates obtained by running the model in equation 1.5. Outcome mean, low-flyer: Effective Leader = 0.41 ; Team Effectiveness = 0.34; Job Satisfaction = 0.35; Autonomy = 0.32; Company effectiveness = 0.34.

Figure 1.66: Heterogeneous effects of high-flyer managers by whether worker changes job

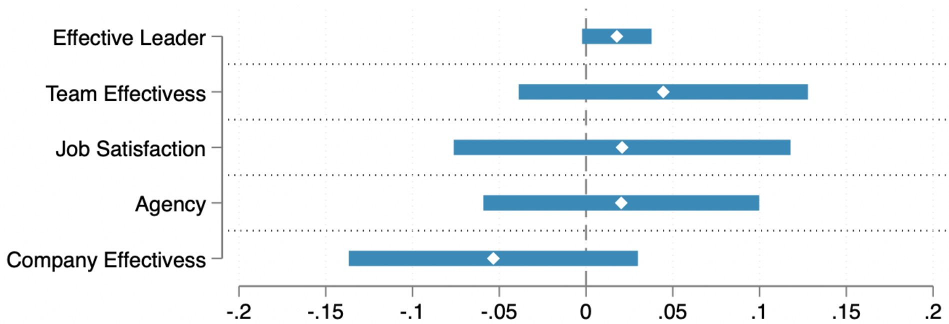

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Survey indices are the first principal components of various survey questions, grouped together by theme as detailed in Appendix Table 1.24. I use binary variables: probability of answering 5 out of 5-point Likert Scale. Estimates obtained by running the model in equation 1.5 interacting indicator for high-flyer manager with an indicator for whether the worker changes job.

Figure 1.67: High-flyer manager and survey evidence, robustness to only first year since manager transition

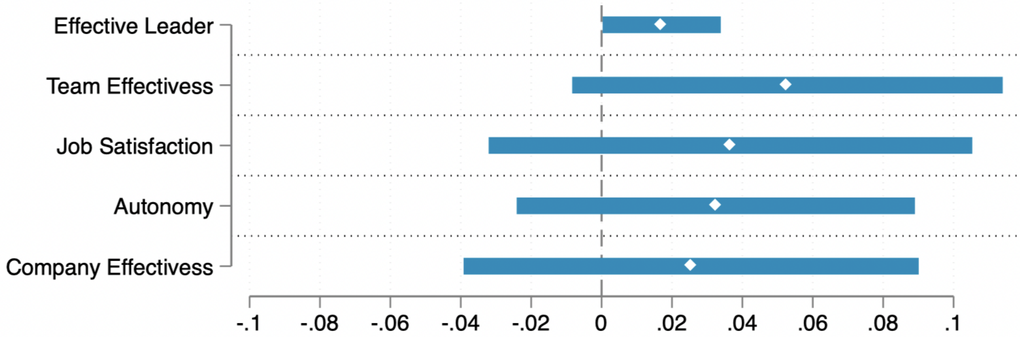

Notes. An observation is a worker-year-month. 90% confidence intervals used and standard errors clustered by manager. Survey indices are the first principal components of various survey questions, grouped together by theme as detailed in Appendix Table 1.24. I use binary variables: probability of answering 5 out of 5-point Likert Scale. Estimates obtained by running the model in equation 1.5 and, unlike Figure 1.23, by selecting only the first year since the manager transition.

Figure 1.68: Share of high-flyers by factory

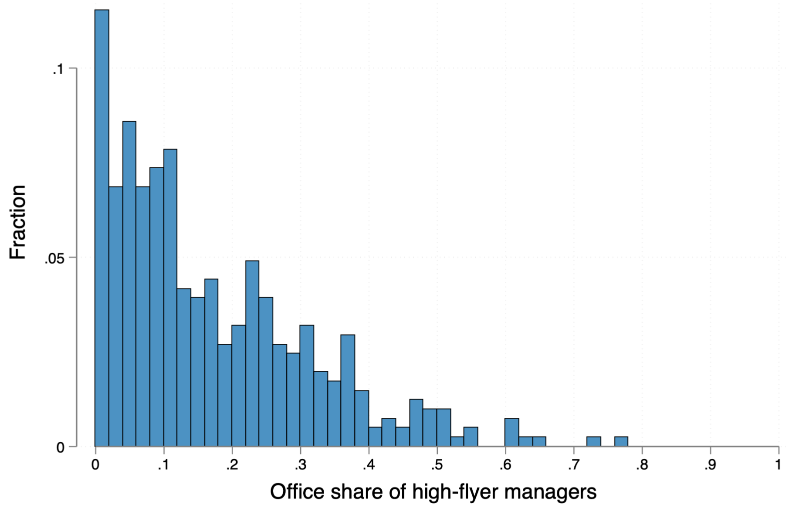

Notes. This figure shows the distribution of high-flyer managers across factories.

Figure 1.69: Factory productivity and share of high-flyer managers

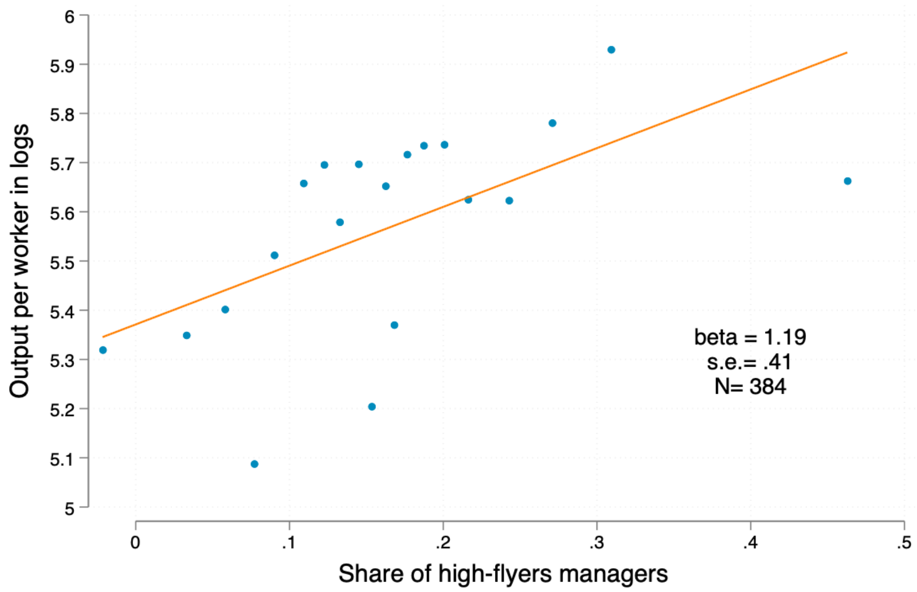

Notes. An observation is a factory-year. Standard errors clustered by factory-year. The y-axis is output per worker in logs and the x-axis is the share of high-flyers in the establishment. Controls include: country, product category and year fixed effects, share of managers, number of blue-collar and white-collar workers.

Figure 1.70: Effects of gaining a high-flyer manager on the probability of making at least one lateral transfer

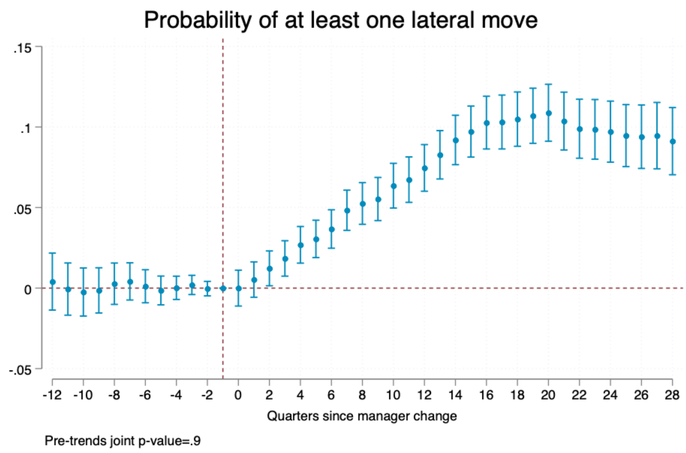

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager.

Figure 1.71: Effects of gaining a high-flyer manager on lateral transfers, poisson model

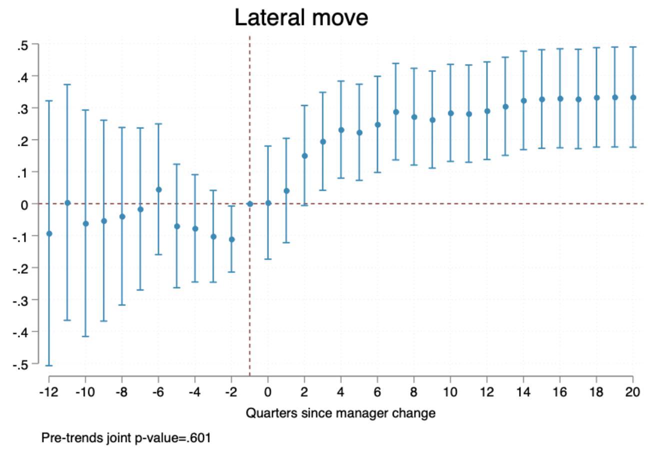

Notes. An observation is a worker-year-month. Reporting the exponentiated coefficients (incidence rate ratios). All coefficients are estimated from a single regression as in equation 1.3 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager.

Figure 1.72: Effects of gaining a high-flyer manager on work-level promotions, poisson model

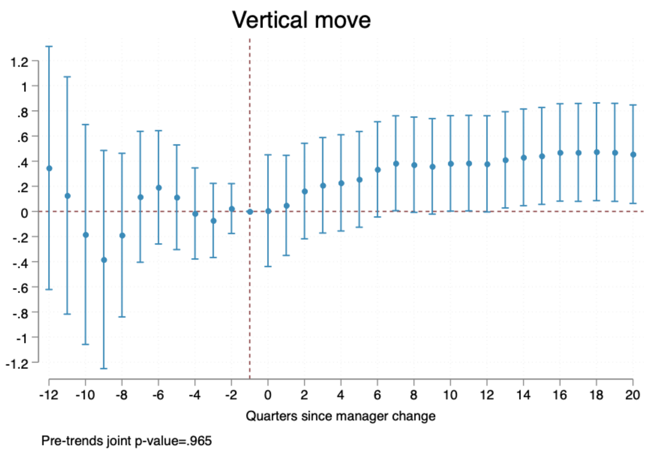

Notes. An observation is a worker-year-month. Reporting the exponentiated coefficients (incidence rate ratios). All coefficients are estimated from a single regression as in equation 1.3 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager.

Figure 1.73: Effects of gaining a high-flyer manager on salary grade increases, poisson model

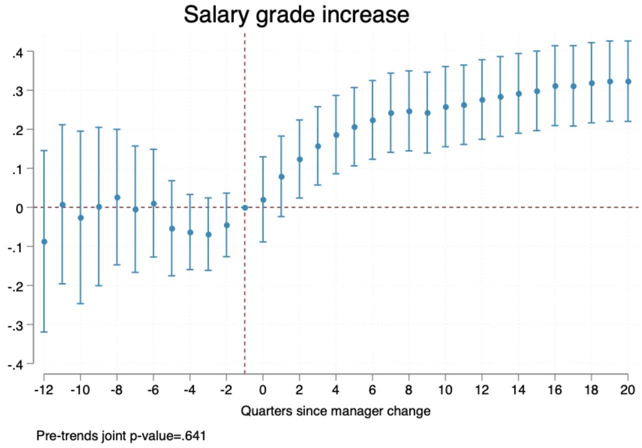

Notes. An observation is a worker-year-month. Reporting the exponentiated coefficients (incidence rate ratios). All coefficients are estimated from a single regression as in equation 1.3 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager.

Figure 1.74: Effects of gaining a high-flyer manager on lateral transfers, single cohort

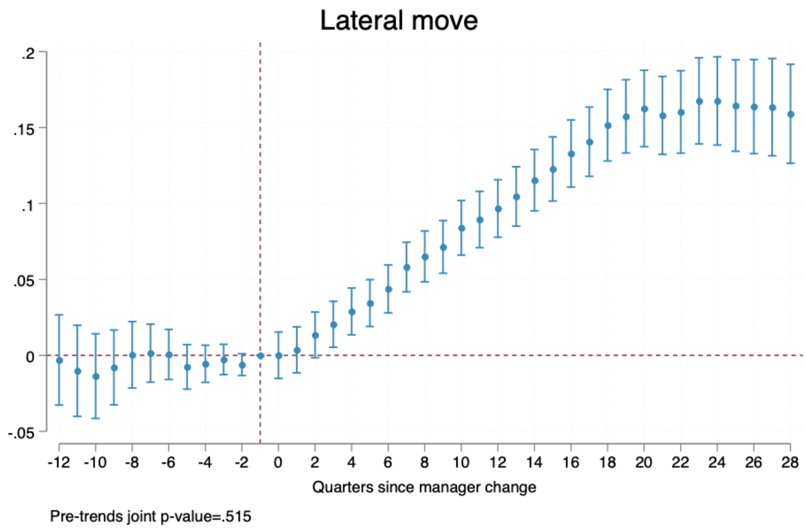

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The workers who experience an event are restricted to those who have it before January 2015.

Figure 1.75: Effects of gaining a high-flyer manager on cross-functional transfers, single cohort

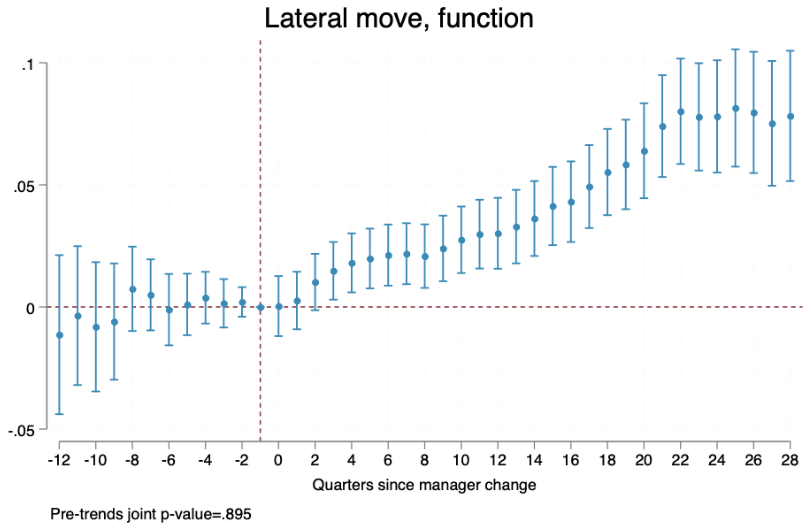

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The workers who experience an event are restricted to those who have it before January 2015.

Figure 1.76: Effects of gaining a high-flyer manager on work-level promotions, single cohort

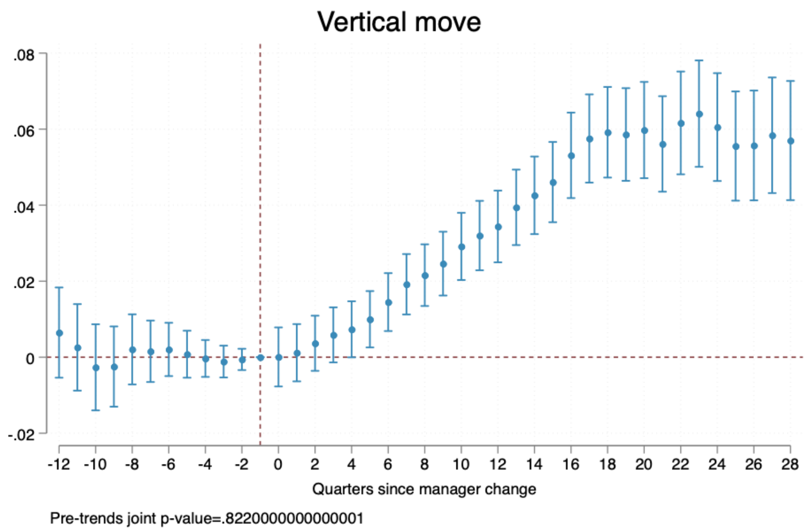

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The workers who experience an event are restricted to those who have it before January 2015.

Figure 1.77: Effects of gaining a high-flyer manager on exit, single cohort

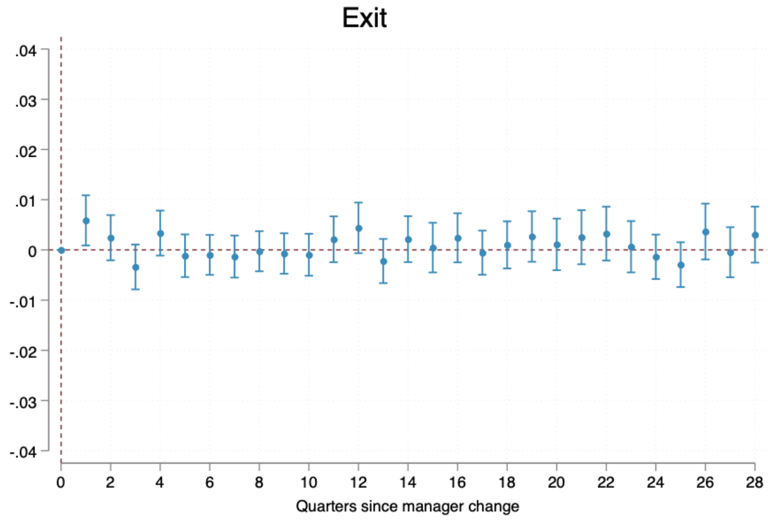

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The workers who experience an event are restricted to those who have it before January 2015.

Figure 1.78: Effects of gaining a high-flyer manager on salary grade increases, single cohort

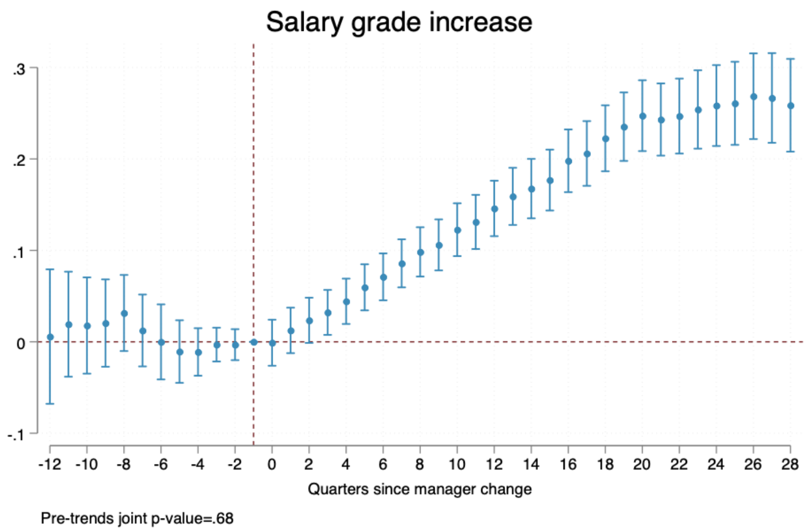

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The workers who experience an event are restricted to those who have it before January 2015.

Figure 1.79: Effects of gaining a high-flyer manager on lateral transfers, new hires

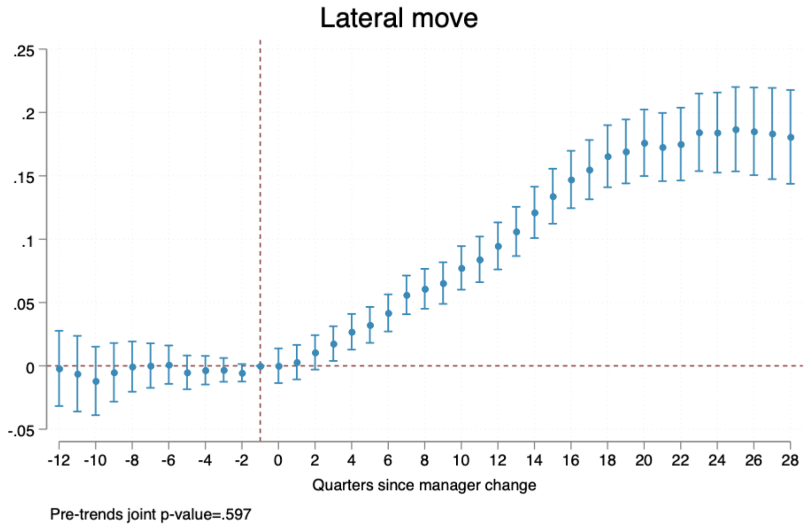

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. Sample restricted to new hires only (with strictly less than one year of tenure).

Figure 1.80: Effects of gaining a high-flyer manager on cross-functional transfers, new hires

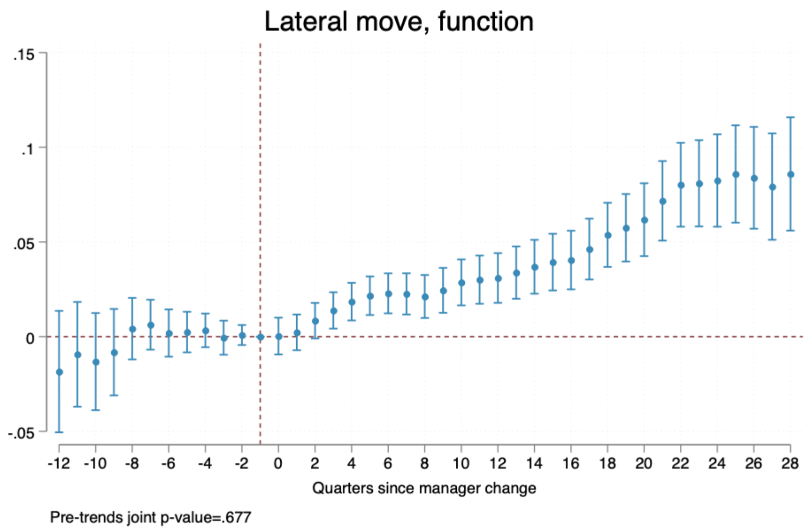

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. Sample restricted to new hires only (with strictly less than one year of tenure).

Figure 1.81: Effects of gaining a high-flyer manager on work-level promotions, new hires

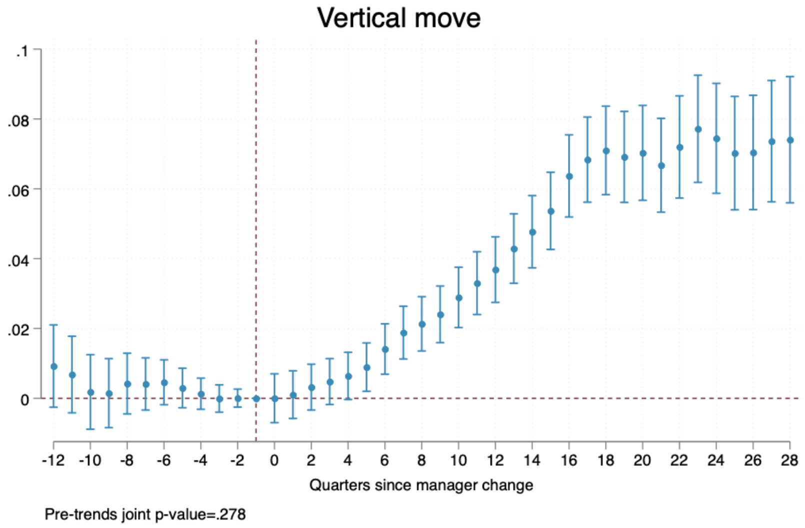

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. Sample restricted to new hires only (with strictly less than one year of tenure).

Figure 1.82: Effects of gaining a high-flyer manager on exit, new hires

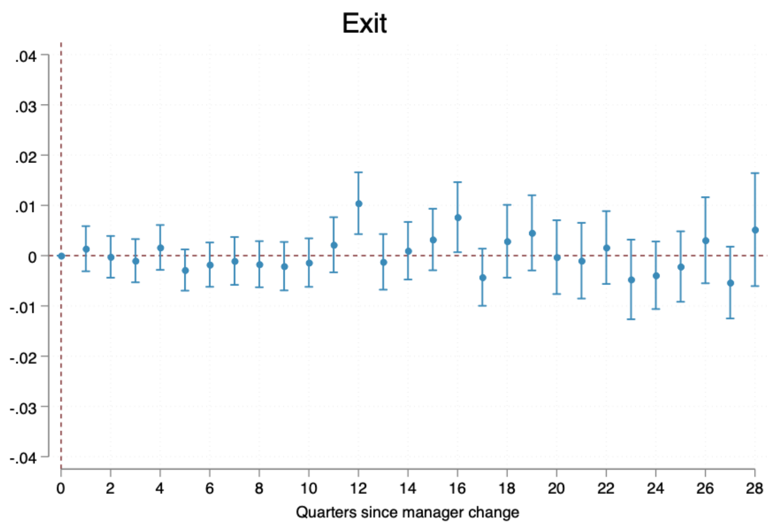

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. Sample restricted to new hires only (with strictly less than one year of tenure).

Figure 1.83: Effects of gaining a high-flyer manager on salary grade increases, new hires

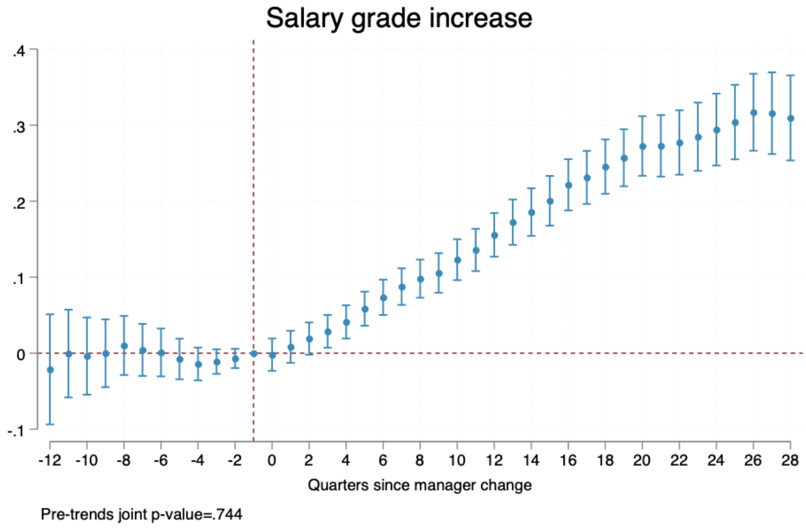

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. Sample restricted to new hires only (with strictly less than one year of tenure).

Figure 1.84: Placebo: lateral transfer

(a) Gain odd-number manager ($\hat{\beta}_{EtoO,s} - \hat{\beta}_{EtoE,s}$) **(b)** Lose odd-number manager ($\hat{\beta}_{OtoE,s} - \hat{\beta}_{OtoO,s}$)

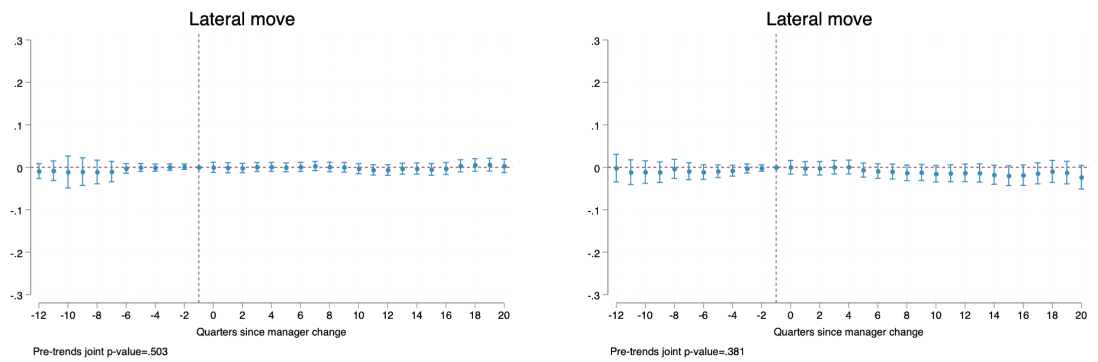

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The scale is the same as the largest of the scales in the corresponding graphs that use the high-flyer manager definition.

Figure 1.85: Placebo: work-level promotion

(a) Gain odd-number manager ($\hat{\beta}_{EtoO,s} - \hat{\beta}_{EtoE,s}$) **(b)** Lose odd-number manager ($\hat{\beta}_{OtoE,s} - \hat{\beta}_{OtoO,s}$)

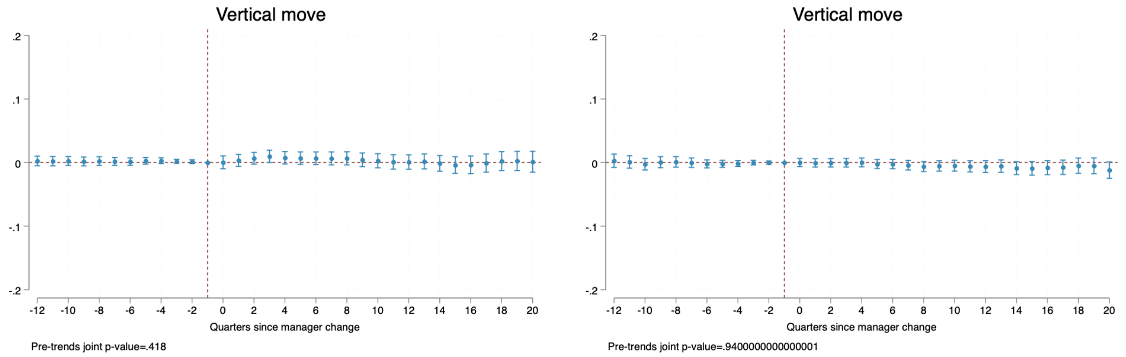

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The scale is the same as the largest of the scales in the corresponding graphs that use the high-flyer manager definition.

Figure 1.86: Placebo: exit from the firm

(a) Gain odd-number manager ($\hat{\beta}_{EtoO,s} - \hat{\beta}_{EtoE,s}$) **(b)** Lose odd-number manager ($\hat{\beta}_{OtoE,s} - \hat{\beta}_{OtoO,s}$)

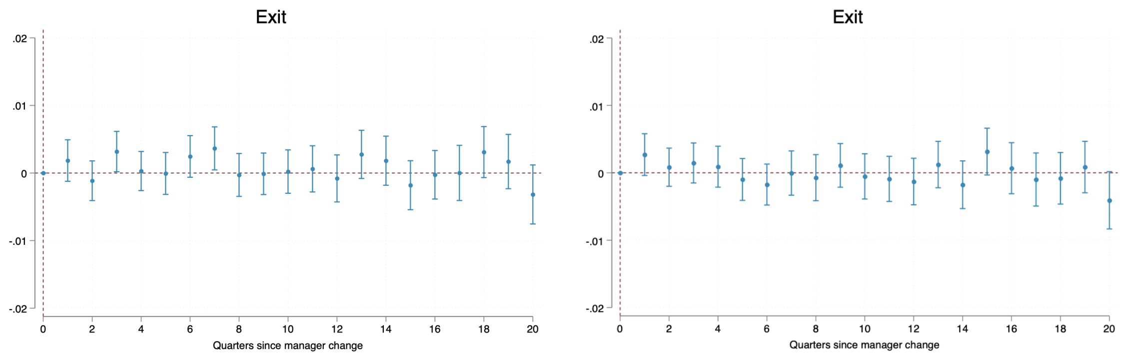

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The scale is the same as the largest of the scales in the corresponding graphs that use the high-flyer manager definition.

Figure 1.87: Placebo: salary grade increase

(a) Gain odd-number manager ($\hat{\beta}_{EtoO,s} - \hat{\beta}_{EtoE,s}$) **(b)** Lose odd-number manager ($\hat{\beta}_{OtoE,s} - \hat{\beta}_{OtoO,s}$)

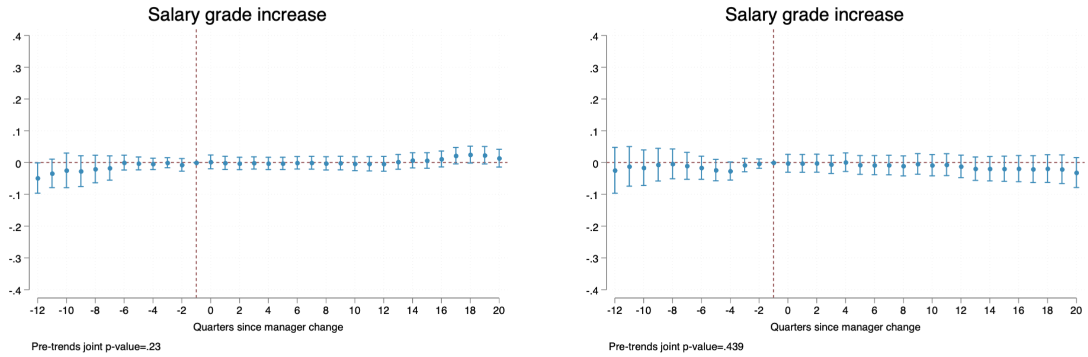

Notes. An observation is a worker-year-month. All coefficients are estimated from a single regression as in equation 1.2 and are aggregated to the quarterly level for ease of presentation. 95% confidence intervals used and standard errors clustered by manager. The scale of each graph is the same as the corresponding graph that uses the high-flyer manager definition.

Figure 1.88: Probability of having a high-flyer manager following the manager of the transition

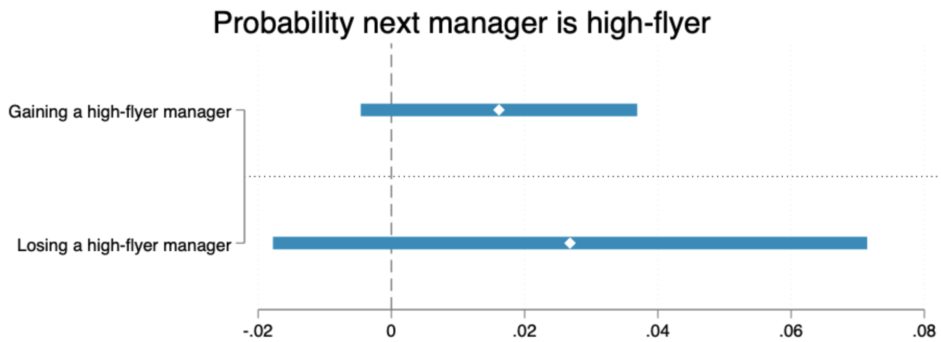

Notes. An observation is a worker. This figure shows the probability for the worker to have a high-flyer manager after the manager transition, separately by the type of manager transition. 90% confidence intervals used and standard errors clustered by manager.

Figure 1.89: Graduate scheme: pay gap by exposure to high-flyer manager during the program

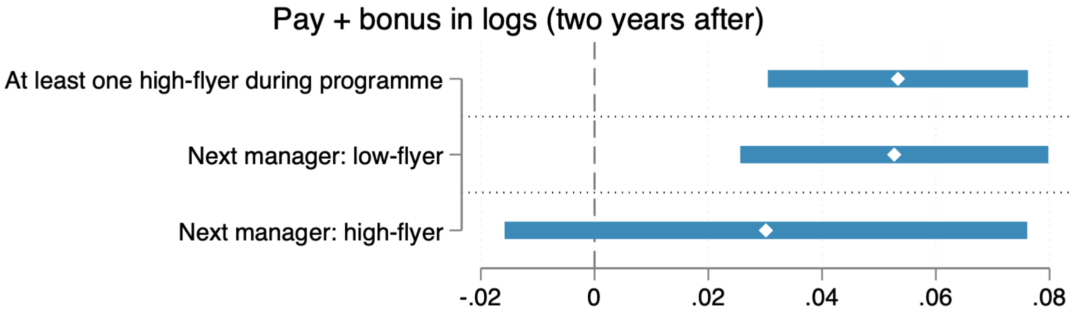

Notes. An observation is a worker-month. 90% confidence intervals used and standard errors clustered by manager. This figure shows the pay differences between workers by whether the worker was exposed to at least one high-flyer manager during the graduate scheme (an entry-level 3-year program into the company for college graduates), controlling for country and year-month fixed effects. At the end of the program, the worker is either allocated to a permanent job in the company under a standard contract or is dismissed. The outcome is total salary in the first two years after the program. The first row looks at all trainees; the second (third) row only considers trainees that move to a low (high)-flyer manager because of the managers' rotation policy.

1.13 Appendix Tables

Table 1.13: Function transition matrix, Low to High Manager Change

Low to High, Pre vs. Post	Customer Development	Finance	General Management	Human Resources	Information Technology	Marketing	Research/Development	Supply Chain	Other	Total
Customer Development	91.64	0.32	0.71	0.39	0.00	6.07	0.00	0.63	0.24	100.00
Finance	1.35	89.68	0.68	0.17	0.51	1.02	0.51	4.57	1.52	100.00
General Management	7.06	4.71	68.24	3.53	3.53	9.41	1.18	1.18	1.18	100.00
Human Resources	2.01	0.50	1.51	90.45	0.00	1.51	0.00	1.01	3.02	100.00
Information Technology	0.00	1.03	0.00	0.00	88.66	2.06	0.00	1.03	7.22	100.00
Marketing	8.07	0.40	1.72	0.26	0.13	88.36	0.26	0.26	0.53	100.00
Research/Development	0.00	0.00	0.00	0.00	0.00	0.70	94.41	4.20	0.70	100.00
Supply Chain	1.05	0.76	0.19	0.19	0.29	0.38	0.67	95.52	0.95	100.00
Other	2.41	3.61	1.20	0.00	1.20	1.20	1.20	1.20	87.95	100.00
Total	29.35	12.97	2.11	4.52	2.27	18.02	3.49	24.60	2.67	100.00

Notes. Biggest eight functions only (98% of employment). Rows indicate the functions at the start while columns indicate the functions 36 months after the manager transition.

Table 1.14: Function transition matrix, Low to Low Manager Change

Low to Low, Pre vs. Post	Customer Development	Finance	General Management	Human Resources	Information Technology	Marketing	Research/Development	Supply Chain	Other	Total
Customer Development	94.84	0.20	0.36	0.16	0.07	3.44	0.02	0.72	0.20	100.00
Finance	2.75	88.22	0.51	0.32	0.90	0.90	0.06	3.97	2.37	100.00
General Management	9.02	17.29	27.82	2.26	22.56	9.77	1.50	7.52	2.26	100.00
Human Resources	1.67	0.48	0.24	91.39	0.00	0.96	0.24	4.55	0.48	100.00
Information Technology	0.40	0.92	0.26	0.13	93.14	0.53	0.00	1.45	3.17	100.00
Marketing	10.54	0.41	0.87	0.12	0.06	86.26	0.29	0.82	0.64	100.00
Research/Development	0.26	0.04	0.00	0.04	0.00	0.68	95.74	3.05	0.19	100.00
Supply Chain	1.60	0.59	0.09	0.24	0.37	0.40	0.79	95.52	0.40	100.00
Other	1.30	4.56	1.30	0.98	7.49	3.91	1.63	7.49	71.34	100.00
Total	31.71	8.32	0.52	2.36	4.50	9.94	14.68	26.10	1.87	100.00

Notes. Biggest eight functions only (98% of employment). Rows indicate the functions at the start while columns indicate the functions 36 months after the manager transition.

Table 1.15: Function transition matrix, High to Low Manager Change

High to Low, Pre vs. Post	Customer Development	Finance	General Management	Human Resources	Information Technology	Marketing	Research/Development	Supply Chain	Other	Total
Customer Development	90.62	0.33	0.55	0.22	0.00	6.84	0.00	1.21	0.22	100.00
Finance	1.60	87.77	0.80	0.27	0.80	1.33	0.00	5.59	1.86	100.00
General Management	21.62	2.70	37.84	0.00	0.00	35.14	0.00	2.70	0.00	100.00
Human Resources	0.68	0.00	0.00	92.57	0.00	2.70	0.00	3.38	0.68	100.00
Information Technology	1.87	0.93	0.93	0.00	94.39	0.00	0.00	0.93	0.93	100.00
Marketing	7.61	0.34	1.69	0.85	0.00	87.14	0.68	1.35	0.34	100.00
Research/Development	0.00	0.00	0.00	0.00	0.98	0.98	94.12	3.92	0.00	100.00
Supply Chain	1.79	0.74	0.15	0.30	0.45	0.45	0.45	95.09	0.60	100.00
Other	2.44	1.22	1.22	2.44	8.54	1.22	1.22	0.00	81.71	100.00
Total	29.69	11.35	1.16	4.93	3.81	19.99	3.44	22.84	2.78	100.00

Notes. Biggest eight functions only (98% of employment). Rows indicate the functions at the start while columns indicate the functions 36 months after the manager transition.

Table 1.16: Function transition matrix, High to High Manager Change

High to High, Pre vs. Post	Customer Development	Finance	General Management	Human Resources	Information Technology	Marketing	Research/Development	Supply Chain	Other	Total
Customer Development	90.03	0.47	1.27	0.47	0.00	7.12	0.16	0.32	0.16	100.00
Finance	1.39	87.50	0.46	0.00	0.46	2.31	0.00	6.48	1.39	100.00
General Management	8.57	0.00	71.43	2.86	0.00	11.43	0.00	2.86	2.86	100.00
Human Resources	0.00	0.00	1.33	94.00	0.67	0.67	0.00	1.33	2.00	100.00
Information Technology	0.00	0.00	0.00	0.00	95.08	1.64	0.00	1.64	1.64	100.00
Marketing	7.21	0.19	2.34	0.19	0.00	88.30	0.00	0.78	0.97	100.00
Research/Development	0.00	0.00	0.00	0.00	0.00	0.00	100.00	0.00	0.00	100.00
Supply Chain	1.56	1.17	0.78	0.39	0.39	2.73	1.56	91.41	0.00	100.00
Other	0.00	2.33	6.98	4.65	0.00	6.98	0.00	0.00	79.07	100.00
Total	32.12	10.27	2.76	7.77	3.18	27.06	0.89	13.45	2.50	100.00

Notes. Biggest eight functions only (98% of employment). Rows indicate the functions at the start while columns indicate the functions 36 months after the manager transition.

Table 1.17: Sub-function transition matrix for Supply Chain, Low to High Manager Change

Low to High, Pre vs. Post	Demand Planning	Engineering	Logistics	Make	Planning	Procurement	Quality	SC Customer Service	SC General Mgmt	Other, Same Func.	Other, Diff. Func.	Total
Demand Planning	40.58	0.00	10.14	7.25	23.19	4.35	0.00	2.90	2.90	1.45	7.25	100.00
Engineering	0.00	78.26	0.00	15.94	0.00	0.00	1.45	0.00	2.90	0.00	1.45	100.00
Logistics	1.09	1.09	48.91	3.26	13.04	4.35	1.09	19.57	2.17	0.00	5.43	100.00
Make	0.38	16.92	1.88	59.77	7.89	0.38	4.89	0.00	2.26	3.38	2.26	100.00
Planning	7.34	1.69	7.91	3.39	61.58	1.69	0.56	5.08	3.95	1.13	5.65	100.00
Procurement	0.76	0.00	1.52	2.27	3.03	87.12	0.76	1.52	0.76	0.00	2.27	100.00
Quality	0.00	2.50	0.00	12.50	5.00	2.50	75.00	0.00	0.00	2.50	0.00	100.00
SC Customer Service	10.39	0.00	7.79	2.60	11.69	0.00	0.00	58.44	1.30	0.00	7.79	100.00
SC General Mgmt	9.78	4.35	3.26	2.17	13.04	9.78	2.17	5.43	35.87	2.17	11.96	100.00
Other, Same Func.	0.00	5.56	2.78	8.33	11.11	2.78	11.11	0.00	2.78	55.56	0.00	100.00
Other, Diff. Func.	0.00	0.03	0.25	0.37	0.16	0.16	0.03	0.22	0.28	0.00	98.51	100.00
Total	1.43	2.60	2.13	4.94	4.54	3.32	1.26	2.06	1.50	0.82	75.40	100.00

Notes. Supply Chain (SC) only (25% of employment). Rows indicate the functions at the start while columns indicate the functions 36 months after the manager transition.

Table 1.18: Sub-function transition matrix for Supply Chain, Low to Low Manager Change

Low to Low, Pre vs. Post	Demand Planning	Engineering	Logistics	Make	Planning	Procurement	Quality	SC Customer Service	SC General Mgmt	Other, Same Func.	Other, Diff. Func.	Total
Demand Planning	51.91	0.85	2.13	2.13	19.15	2.55	0.43	7.66	6.81	0.43	5.96	100.00
Engineering	0.88	80.31	0.44	10.72	0.66	1.31	1.75	0.44	0.66	1.53	1.31	100.00
Logistics	4.67	0.85	57.75	1.70	6.58	2.76	0.21	16.77	3.18	0.21	5.31	100.00
Make	0.87	17.52	3.38	46.83	5.98	2.17	8.93	0.95	6.33	5.64	1.39	100.00
Planning	10.90	0.47	3.74	2.96	63.08	3.43	0.93	4.98	3.43	1.25	4.83	100.00
Procurement	0.68	1.02	2.20	0.68	2.20	82.06	0.68	0.85	2.20	0.85	6.60	100.00
Quality	0.38	3.07	0.77	8.81	3.07	1.53	77.39	0.00	0.00	2.68	2.30	100.00
SC Customer Service	5.96	0.00	4.64	0.33	9.60	1.32	0.66	65.56	1.66	0.33	9.93	100.00
SC General Mgmt	5.54	7.27	6.57	6.92	9.34	3.11	2.77	5.19	36.68	5.88	10.73	100.00
Other, Same Func.	0.65	6.54	1.96	16.34	1.96	1.31	3.27	0.00	1.31	62.75	3.92	100.00
Other, Diff. Func.	0.10	0.12	0.21	0.15	0.24	0.29	0.16	0.26	0.36	0.09	98.02	100.00
Total	1.59	3.62	2.38	4.04	3.77	3.48	2.04	2.23	1.71	1.25	73.90	100.00

Notes. Supply Chain (SC) only (25% of employment). Rows indicate the functions at the start while columns indicate the functions 36 months after the manager transition.

Table 1.19: Sub-function transition matrix for R&D, Low to High Manager Change

Low to High, Pre vs. Post	Consumer Technical Insight	Packaging Development	Processing Development	Product Development	R&D General Mgmt	Science&Tech Discovery	Other, Same Func.	Other, Diff. Func.	Total
Consumer Technical Insight	84.21	0.00	0.00	5.26	5.26	0.00	0.00	5.26	100.00
Packaging Development	0.00	100.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00
Processing Development	6.25	6.25	50.00	18.75	0.00	6.25	6.25	6.25	100.00
Product Development	3.85	1.92	0.00	75.00	5.77	5.77	0.00	7.69	100.00
R&D General Mgmt	4.55	0.00	0.00	13.64	72.73	4.55	0.00	4.55	100.00
Science&Tech Discovery	0.00	0.00	7.69	7.69	0.00	76.92	0.00	7.69	100.00
Other, Same Func.	0.00	0.00	0.00	0.00	0.00	0.00	100.00	0.00	100.00
Other, Diff. Func.	0.02	0.07	0.10	0.02	0.07	0.00	0.05	99.66	100.00
Total	0.49	0.40	0.30	1.12	0.54	0.35	0.28	96.51	100.00

Notes. R&D only (8% of employment). Rows indicate the functions at the start while columns indicate the functions 36 months after the manager transition.

Table 1.20: Sub-function transition matrix for R&D, Low to Low Manager Change

Low to Low, Pre vs. Post	Consumer Technical Insight	Packaging Development	Processing Development	Product Development	R&D General Mgmt	Science&Tech Discovery	Other, Same Func.	Other, Diff. Func.	Total
Consumer Technical Insight	75.34	0.68	0.00	7.53	4.79	4.79	2.05	4.79	100.00
Packaging Development	0.58	84.88	1.16	4.65	2.03	0.00	1.16	5.52	100.00
Processing Development	0.00	0.47	75.81	13.49	3.26	1.86	0.47	4.65	100.00
Product Development	2.62	2.40	3.72	80.00	2.84	0.98	2.40	5.03	100.00
R&D General Mgmt	1.33	1.67	2.00	17.33	55.67	2.33	14.00	5.67	100.00
Science&Tech Discovery	0.80	0.00	0.40	3.42	1.01	91.35	2.62	0.40	100.00
Other, Same Func.	0.84	0.00	0.42	2.53	5.91	0.84	84.39	5.06	100.00
Other, Diff. Func.	0.01	0.07	0.03	0.13	0.07	0.01	0.03	99.66	100.00
Total	0.83	1.87	1.22	5.00	1.38	2.74	1.64	85.32	100.00

Notes. R&D only (8% of employment). Rows indicate the functions at the start while columns indicate the functions 36 months after the manager transition.

1.14 Theoretical Appendix

In this section, I provide additional details for the predictions in sub-section 1.7.4 and the discussion in sub-section 1.7.5.

The expected worker productivities used to derive Prediction 1 and 2 are illustrated in Table 1.21.

Table 1.21: Framework: expected productivity matrix by initial job allocation

		Manager type	
		Good	*Bad*
	Social $\longrightarrow^{?}$ *Analytical*	m^A	$m^A + \epsilon_1^A$
Current job$_1$	*Social* $\longrightarrow^{?}$ *Social*	$\tau + m^S$	$1 + m^S + \epsilon_1^S$
$\longrightarrow^{?}$ Next job$_2$	*Analytical* $\longrightarrow^{?}$ *Analytical*	$\tau + m^A$	$1 + m^A + \epsilon_1^A$
	Analytical $\longrightarrow^{?}$ *Social*	m^S	$m^S + \epsilon_1^S$

Notes. This table shows the expected worker productivity computed by the manager.
It depends on the worker job move and the manager type.
The worker starts with no experience in either the analytical or the social job.
The worker can move to either the analytical or the social job.

In sub-section 1.7.5 I discuss the conditions under which: (i) moving from a bad to a good manager compared to moving from a bad to another bad manager (*gaining a good manager*) leads to higher job transfer rates and future productivity, and (ii) moving from a good to a bad manager compared to moving from a good to another good manager (*losing a good manager*) has no differential impact on job transfer rates and future productivity. This requires me to step outside the one period set-up and evaluate the equilibrium path for two periods. I use the worker expected productivities illustrated in Table 1.22 and Table 1.23.

First, consider the effects of losing a good manager. As the first manager is good, the probability that the worker is in the bad job match (the social job given the model set-up) is zero, given Prediction 1 (in sub-section 1.7.4).

A good manager never moves the worker. A bad manager never moves the worker if she knows that the previous manager of the worker was good. Average future worker productivity will be the same among the two manager types if $\tau = 1$ (no difference in teaching between a good and bad manager) or if there are decreasing returns to learning-by-doing. In particular, the accumulation of experience goes to zero after one period on the job. Although this results in a coarse restriction to the evolution of learning-by-doing (which is a consequence of the simple model set-up), it is plausible that learning exhibits decreasing returns.

Second, consider the effects of gaining a good manager. As the first manager is bad, there is a non-zero probability of the worker being in the bad job-match (i.e. the social job).

If the worker is in the social job, a good manager moves her with probability 1 to the analytical job if $m^a - m^s > 2$ (job allocation is more important than learning-by-doing). On the other hand, a bad manager moves her with probability $\varPhi\left(\frac{(m^A-m^S)-2}{\sigma_\epsilon^{2A}+\sigma_\epsilon^{2S}}\right) < 1$ (if first job was social) or with probability $\varPhi\left(\frac{(m^A-m^S)}{\sigma_\epsilon^{2A}+\sigma_\epsilon^{2S}}\right) < 1$ (if first job was analytical). If the worker is in the analytical job, a good manager never moves the worker to a social job, while a bad manager moves her to the social job with probability $1-\varPhi\left(\frac{m^A-m^S}{\sigma_\epsilon^{2A}+\sigma_\epsilon^{2S}}\right)$ (if first job was social) or probability $1-\varPhi\left(\frac{m^A-m^S+2}{\sigma_\epsilon^{2A}+\sigma_\epsilon^{2S}}\right)$ (if first job was analytical).

Note that both $\left(\varPhi\left(\frac{(m^A-m^S)-2}{\sigma_\epsilon^{2A}+\sigma_\epsilon^{2S}}\right)+1-\varPhi\left(\frac{m^A-m^S}{\sigma_\epsilon^{2A}+\sigma_\epsilon^{2S}}\right)\right)$ (if first job was social) and $\left(\varPhi\left(\frac{(m^A-m^S)}{\sigma_\epsilon^{2A}+\sigma_\epsilon^{2S}}\right)+1-\varPhi\left(\frac{m^A-m^S+2}{\sigma_\epsilon^{2A}+\sigma_\epsilon^{2S}}\right)\right)$ (if first job was analytical) are less than one. Hence, there is a higher chance of the worker changing job when the second manager is good compared to when it is bad. It follows that average future productivity is also higher as the worker is more likely to end up to the right job match with a good manager.

Table 1.22: Framework: expected productivity by manager transition, first job is analytical

		Manager transition			
		$Bad_1, Good_2$	Bad_1, Bad_2	$Good_1, Bad_2$	$Good_1, Good_2$
Job_1	$Anal._1 Anal.^?_2 Anal._3$	$1+\tau+m^A$	$2+m^A+\epsilon_2^A$	$\tau+1+m^A+\epsilon_2^A$	$2\tau+m^A$
$= Anal.$	$Anal._1 Social^?_2 Anal._3$	$1+m^A$	$1+m^A+\epsilon_2^A$	$\tau+m^A+\epsilon_2^A$	$\tau+m^A$
Job_2	$Anal._1 Social^?_2 Social_3$	$\tau+m^S$	$1+m^S+\epsilon_2^S$	$1+m^S+\epsilon_2^S$	$\tau+m^S$
$^?$ Job_3	$Anal._1 Anal.^?_2 Social_3$	m^S	$m^S+\epsilon_2^S$	$m^S+\epsilon_2^S$	m^S

Notes. This table shows the expected worker productivity computed by the manager.
It depends on the worker history in terms of jobs and manager types.
The worker starts with no experience in the analytical job.
The worker can move to either the social or the analytical job in period 2 and 3.

Table 1.23: Framework: expected productivity by manager transition, first job is social

		Manager transition			
		$Bad_1, Good_2$	Bad_1, Bad_2	$Good_1, Bad_2$	$Good_1, Good_2$
Job_1	$Social_1 Anal.^?_2 Anal._3$	$\tau + m^A$	$1 + m^A + \epsilon_2^A$	$1 + m^A + \epsilon_2^A$	$\tau + m^A$
$= Social.$	$Social_1 Social^?_2 Anal._3$	m^A	$m^A + \epsilon_2^A$	$m^A + \epsilon_2^A$	m^A
Job_2	$Social_1 Social^?_2 Social_3$	$1 + \tau + m^S$	$2 + m^S + \epsilon_2^S$	$1 + \tau + m^S + \epsilon_2^S$	$2\tau + m^S$
$^? \text{ Job}_3$	$Social_1 Anal.^?_2 Social_3$	$1 + m^S$	$1 + m^S + \epsilon_2^S$	$\tau + m^S + \epsilon_2^S$	$\tau + m^S$

Notes. This table shows the expected worker productivity computed by the manager.

It depends on the worker history in terms of jobs and manager types.

The worker starts with no experience in the social job.

The worker can move to either the social or the analytical job in period 2 and 3.

1.15 Data Appendix

1.15.1 Measuring task distance between occupations

Occupations, as discrete classification units, can be viewed as vectors of tasks to be carried out by workers. I manually match the occupation codes in the firm to the Occupational Information Network (O*NET) classification codes and obtain vectors for each occupation o, $q_o^c = (q_{o1}, \ldots, q_{oN})$ where c is skills, activities, abilities, work contexts. These job content measures can be understood as describing a position in the task space. My baseline results make use of the skills vector but they are robust to taking the average of the different vectors. I consider the skills vector as my empirical analogue of the occupation-specific weights on tasks in the conceptual framework: occupations with a high weight in a particular task w, β_o^w, will have a high q_{ow}^{skills}. The Occupational Information Network (O*NET) offer multiple sources for job content descriptors, and has been used frequently in empirical work on job tasks (Autor, 2013).

I follow Gathmann and Schönberg (2010) and define the angular separation between occupation j and occupation k as a measure of similarity using task vectors q_j^{skills} and q_k^{skills}:

$$AngSim_{jk} = \frac{\sum_{s=1}^{S}(q_{sj} \times q_{sk})}{\left((\sum_{s=1}^{S} q_{sj}^2) \times (\sum_{s=1}^{S} q_{sk}^2)\right)^{1/2}}$$

This angular separation measure defines the distance between two occupations as the cosine angle between their positions in vector space. I define $(1 - AngSim_{jk})$ as the distance between occupation j and occupation k: $Dist_{jk} = (1 - AngSim_{jk})$. The measure ranges between zero and one. It is zero for occupations that use identical skill sets and unity if two occupations use completely different skills sets. The measure will be closer to zero the more two occupations overlap in their skill requirements.

Measuring similarity between two vectors by the angular separation has first been proposed by Jaffe (1986) in the innovation literature to characterize the proximity of firm's technologies. Subsequently, a number of other studies have used the measure in various contexts such as spillovers of university research to commercial innovation (Jaffe, 1989), and similarity of tasks performed across occupations (Gathmann and Schönberg (2010)), Cortes and Gallipoli (2018)).

The mean distance between occupations in my data is 0.06, with a standard deviation of 0.14. As the focus here is job moves within the same firm as opposed to moves across firms, there are many moves where task distance is 0, for example between a recruitment specialist and a general talent advisor, both in human resources. The most distant possible move is between a tax administrator in finance and a production supervisor on the factory floor in supply chain.

1.15.2 Survey questions

Every year in September since 2019 all employees are invited via e-mail by the firm's human resources department to provide their perspectives on team dynamics, company strategy, and the overall work environment.[83] On average each year, the survey receives over 45,000 responses, yielding a 55% response rate. Respondents are broadly similar to non-respondents in terms of demographics; they generally tend to be slightly older, higher up in the hierarchy, and are marginally more likely to have a high-flyer manager (Appendix Table 1.25). I only keep respondents who have no missing answers.

Table 1.24 shows the questions and how they were grouped together in the indices used in the analysis (using the loadings on the first principal component).

[83]The survey was launched in 2017, however, in 2017 and 2018, it was only sent to a random sample of 20% of employees.

Table 1.24: Variable construction - Survey Measures

Variable	*Components*	*Possible answers*
Panel A: Team effectivess		
Team inclusive	In my team, we have an inclusive working environment	1 Strongly disagree - 5 Strongly agree
Team agility	I feel that over the last 12months the speed & agility has improved in my teams	1 Strongly disagree - 5 Strongly agree
Trust leaders	I trust the Senior leaders in my part of the organization	1 Strongly disagree - 5 Strongly agree
Leaders strategy	Leaders in my part of the org. clearly demonstrate strategy in their behaviour	1 Strongly disagree - 5 Strongly agree
Leaders inclusive	Leadership in my part of the org. visibly stands for diversity & inclusion	1 Strongly disagree - 5 Strongly agree
Customers at heart	My team puts the needs of our customers at the heart of everything we do	1 Strongly disagree - 5 Strongly agree
Panel B: Worker autonomy		
Focus on performance	I am able to manage distractions and focus on what matters	1 Strongly disagree - 5 Strongly agree
Access learning	I can access the learning resources I need to do my job effectively	1 Strongly disagree - 5 Strongly agree
Prioritization	I have control over prioritising tasks when facing multiple demands at work	1 Strongly disagree - 5 Strongly agree
Development	I am satisfied with my development opportunities at Company	1 Strongly disagree - 5 Strongly agree
Wellbeing	I believe that Company cares about my Wellbeing	1 Strongly disagree - 5 Strongly agree
Report unethical behavior	I feel able to report potential bus. principle breaches w/o fear of retaliation	1 Strongly disagree - 5 Strongly agree
Panel C: Job satisfaction		
Work life balance	I can maintain a reasonable balance between my personal life and work life	1 Strongly disagree - 5 Strongly agree
Job satisfaction	Overall, I am extremely satisfied with Company as a place to work	1 Strongly disagree - 5 Strongly agree
Refer Company	I would gladly refer a friend or family member to Company for employment	1 Strongly disagree - 5 Strongly agree
Proud to be at Company	I am proud to say that I work for Company	1 Strongly disagree - 5 Strongly agree
Live purpose in Company	I believe I can live my purpose in Company	1 Strongly disagree - 5 Strongly agree
Leaving Company	I am not seriously considering leaving Company	1 Strongly disagree - 5 Strongly agree
Extra mile	My job inspires me to go the extra mile	1 Strongly disagree - 5 Strongly agree
Panel D: Company effectiveness		
Strategy to win	Company has the right strategy in place to win	1 Strongly disagree - 5 Strongly agree
Sustainability	My job contributes to the sustainability plan and drives sustainable growth	1 Strongly disagree - 5 Strongly agree
Technology	Company processes & technologies available to me make it easier to do my job	1 Strongly disagree - 5 Strongly agree
Competition	Company better than competition at responding rapidly to changes in the market	1 Strongly disagree - 5 Strongly agree
Removing barriers between teams	Company helps me to work efficiently by removing barriers between teams	1 Strongly disagree - 5 Strongly agree
Integrity	I believe that in Company business is conducted with integrity	1 Strongly disagree - 5 Strongly agree
Recommend products	I would recommend Company's products to my family and friends	1 Strongly disagree - 5 Strongly agree
Panel E: Effective line manager		
Effective manager	My line manager is an effective leader	1 Strongly disagree - 5 Strongly agree

Table 1.25: Comparison of non-respondents to respondents - employee annual survey

	(1)	(2)	(3)
	Mean / (SE)		**Difference in means / (p-value)**
Variable	Non-respondents	Survey respondents	Difference
Female	0.432	0.466	0.009***
	(0.495)	(0.499)	(0.000)
Share in Cohort 18-29	0.268	0.206	-0.050***
	(0.443)	(0.404)	(0.000)
Share in Cohort 30-39	0.387	0.404	-0.005**
	(0.487)	(0.491)	(0.017)
Share in Cohort 40-49	0.221	0.247	0.032***
	(0.415)	(0.431)	(0.000)
Share in Cohort 50+	0.124	0.143	0.023***
	(0.330)	(0.350)	(0.000)
Econ, Business, and Admin	0.476	0.488	0.003
	(0.499)	(0.500)	(0.570)
Sci, Engin, Math, and Stat	0.309	0.300	0.007
	(0.462)	(0.458)	(0.112)
Social Sciences and Humanities	0.146	0.147	-0.003
	(0.353)	(0.354)	(0.430)
Other Educ	0.075	0.071	-0.008***
	(0.263)	(0.256)	(0.002)
Tenure (years)	8.199	9.341	1.677***
	(8.765)	(8.937)	(0.000)
Share in Work-level 1	0.819	0.742	-0.105***
	(0.385)	(0.438)	(0.000)
Share in Work-level 2	0.146	0.206	0.079***
	(0.353)	(0.405)	(0.000)
Share in Work-level 3+	0.035	0.052	0.026***
	(0.184)	(0.222)	(0.000)
High-flyer manager	0.127	0.201	0.042***
	(0.333)	(0.401)	(0.000)
Observations	678,557	158,829	837,386

Notes. This table compares average characteristics of the non-respondents (Column 1) to the subset of employees who responded to the employee survey (Column 2). Standard errors clustered at the worker level used. Controlling for office year fixed effects.

Chapter 2

Union Leaders: Experimental Evidence From Myanmar

This chapter is jointly co-authored with Laura Boudreau, Rocco Macchiavello, and Mari Tanaka.

2.1 Introduction

Social movements have been critical drivers of many institutional changes: In the 19th century, the eight-hour day movement, in the early 1900s, the suffragettes, in the 1950s, the civil rights movements, and in this century, the green movement (Della Porta and Diani, 2020), to name but a few. To succeed, social movements must coordinate their members' views and collective actions. Coordinating views requires building *consensus* around common objectives and tactics among diverse members. Coordinating actions requires *mobilizing* members to participate in activities that are characterized by having high, private costs and uncertain, public benefits (Ganz, 2010). But unlike in more formal organizational settings, monetary incentives, contracts, and hierarchies are often unavailable to align views and motivate members.

In the absence of these organizational tools, we hypothesized that *leaders* may play a critical role.[1] Economic theory suggests that leaders may act as coordinators in both group decision-making and mobilization. For example, leaders may build consensus among a group by providing information about the state of the world or payoffs that coordinates views (Hermalin (1998); Caillaud and Tirole (2007); Dewan and Myatt (2008)). They may mobilize group members by communicating that a high-cooperation equilibrium is to be played (Loeper et al., 2014). This paper provides empirical evidence on these theories from experiments on union leaders' roles as coordinators in consensus building and mobilization in Myanmar's burgeoning labor movement

[1] We define leadership in the spirit of Hermalin (2012)'s definition; he states "...one of the essences of leadership is the ability to induce others to follow absent the power to compel or to provide formal contractual incentives...A leader is someone with followers, who follow voluntarily."

– a movement that is broadly representative of the struggles in organizing labor movements in newly industrializing countries (see, e.g., Visser et al. (2019)).

To date, empirical evidence on leaders' roles in group decision-making and in mobilization outside the lab remains scarce due to measurement and identification challenges.[2] On the measurement side, it is difficult to observe many leaders performing the same task. On the identification side, it is difficult to distinguish if a given individual influences others (i.e., is in fact a leader) or if their behavior is simply a more visible emblem of the underlying group dynamics - a version of the well-known "reflection problem" (Manski, 1993).

We overcame these challenges by collaborating with the Confederation of Trade Unions in Myanmar (CTUM), a confederation of labor unions that represented workers' interests in the national minimum wage setting process. In the run-up to the planned May 2020 negotiations, the CTUM organized weekend sessions with workers at 17 garment factories with CTUM-affiliated unions to discuss the minimum wage and to gather systematic information on workers' skills and living costs. We helped the CTUM to organize the discussions and to conduct the surveys, which allowed us to embed multiple experiments to examine (1) whether and how union leaders matter in consensus building regarding the minimum wage level and (2) whether and how they mobilize workers to participate in privately costly activities for the common good. The sessions allowed us to study many different union leaders across many different unions, overcoming the measurement challenge, and to run experiments, overcoming the identification challenge. They also presented a unique opportunity to study leaders' roles in the context of a high-stakes, real-world collective action to influence a policy choice with uncertainty about its success – the CTUM's effort to influence the national minimum-wage level – while avoiding many of the risks associated with mobilization around, for example, protests in the streets.

We present three sets of empirical results. We first document that union leaders are distinct from union members and non-members along key traits that economists identify as relevant for political selection (Caselli and Morelli (2004); Dal Bó et al. (2017)) and that psychologists and organizational sociologists associate with the ability to influence collective outcomes (Judge et al., 2002), respectively. In each factory, the union leadership is structured around an elected union president and executive committee that negotiates with the factory management and coordinates activities with the confederation. Below these formal roles, several (typically) non-elected line leaders (LLs) organize and voice the concerns of other union members. We find that, relative to other workers, presidents have higher Raven Scores. Both presidents and LLs are more altruistic, more extroverted, less neurotic, and more conscientious compared to workers.

[2] Englmaier et al. (2022) is a notable exception; they provide experimental evidence that teams with endogenously-chosen leaders outperform teams without leaders in the context of an escape room challenge.

They have greater grit and locus of control. They also have more work experience.

We then present results from two sets of field experiments. In these experiments, we focus on understanding how LLs influence workers' behavior. There are a few justifications for our choice to focus on LLs. First, LLs are tasked by the union to directly interact, mobilize, and gather and channel the concerns of the workers. Second, as is evident from their traits, LLs resemble formal leaders. There are also a far greater number of them. We can thus observe the behavior of many leaders who, albeit not (yet) formally elected to leadership positions within the union, share many of the traits of (and are likely to subsequently become) union leaders.[3]

In experiment (1) on consensus building, we randomly embedded leaders in group discussions about workers' preferred and expected minimum wage levels. We find that leaders increase workers' consensus around their unions' preferred minimum wage levels by 20%. In terms of how leaders achieve this, groups with leaders are substantially more likely to have a member taking notes and summarizing opinions and have fewer distracted participants, according to field staff observers. Using textual data from discussion transcripts, we find that leaders appear to be introducing influential minimum wage values during the discussion. Following the discussions, workers in groups with leaders self-report higher engagement and perception that the group achieved consensus. Thus, leaders appear to be coordinating both by combining and organizing workers' views and by introducing information. We experimentally varied whether a leader is assigned to a group with workers from her own or from another factory, allowing us to provide evidence that leaders' *social connections* (Bandiera et al., 2009) or their *formal role* (Aghion and Tirole, 1997) in the organization alone cannot explain these results; leaders' own attributes, including their resemblance to formal leaders, matter.

In experiment (2) on mobilization, we invited workers to participate in an unannounced survey on living costs. The survey was a privately costly action, as it required them to sacrifice the remainder of their one weekend day, which conveyed a public benefit, as the CTUM planned to use the data to campaign for its preferred minimum wage level. We induced a strategic complementarity in turnout at the discussion group level by donating to a worker skills training center for each full discussion group that attended the survey. In this way, we mirrored the incentives faced by workers when deciding whether to participate in collective actions such as street demonstrations in support of the CTUM's proposed minimum wage level. We varied whether workers: (i) were invited to the survey by a leader; (ii) were informed about how many discussion group members were invited by a leader; (iii) were told that a leader would observe their decision to participate.

[3]We also conducted an experiment involving presidents. Due to the COVID-19 pandemic, however, we could only cover 60% of the sample. Given the much lower number of presidents (one per factory) compared to LLs, we are underpowered to detect effects in this experiment.

Again, we find that leaders play a coordinating role: Moving from being informed that a leader would invite one group member to being informed that they would invite all but one group member increases take-up by 38%. In contrast, being invited by a leader alone does not increase take-up. Finally, we find suggestive evidence that observation of the workers' choice by a leader increases take-up; our results are consistent with a signaling channel.

This research contributes to three strands of literature. First, it contributes to the literature on leaders' roles in group decision-making and in overcoming collective action problems. A sizable theoretical literature focuses on forms of information provision by leaders that serve to coordinate beliefs or actions (Hermalin (1998); Caillaud and Tirole (2007); Dewan and Myatt (2008); Bolton et al. (2012); Loeper et al. (2014); Akerlof and Holden (2016)). Empirically, the literature is primarily composed of lab experiments (Potters et al. (2007); Komai et al. (2010); Sahin et al. (2015)). More recently, a limited number of field experiments have studied leadership in real-world settings; for example, Grossman and Baldassarri (2012) on sanction enforcement, Jack and Recalde (2015) on signaling and reciprocity, and Englmaier et al. (2022) on the importance of selecting of a leader. We contribute by providing evidence on leaders' roles in group decision-making and in overcoming collective action problems from experiments with many different real-world leaders in a burgeoning labor movement's effort to influence a high-stakes policy-setting process. Further, our experimental designs and rich data sources enable us to provide novel micro-evidence on the mechanisms through which leaders influence outcomes.

Second, we contribute to an emerging empirical literature on the determinants of social movements' formation and growth. One stream of this literature focuses on how information about others' participation affects individuals' decisions to participate in protests; it underscores that coordination problems present an important challenge to turnout and emphasizes the importance of mechanisms to enhance coordination (mainly, communication technology) (González (2020); Enikolopov et al. (2020); Manacorda and Tesei (2020)).[4] A second stream focuses on how the presence of leaders affects individuals' decisions to participate. Dippel and Heblich (2021) and Cagé et al. (2020) provide novel evidence from different historical social movements that exposure to leaders increases participation, but both are constrained in their ability to speak to the mechanisms through which leaders influence collective outcomes. We contribute causal evidence that leaders are an important mechanism to enhance coordination in social movements. We also contribute evidence on the mechanisms through which leaders align group members with the movement's objectives and mobilize them to participate in collective action.

[4]In contrast to other recent papers on this topic, Cantoni et al. (2019) provide causal evidence of strategic substitutability in protest turnout in the context of Hong Kong's long-running democracy movement. In these types of settings, even if leaders do not serve a coordinating role, in principle, they may still play important roles through other channels such as social pressure.

Third, this paper contributes to a growing literature on industrial relations and labor unions in developing countries (Freeman (2010); Tanaka (2020); Boudreau (2021); Macchiavello et al. (2020); Breza et al. (2019); Akerlof et al. (2020); Lin et al. (2019); Corradini et al. (2022)). We contribute well-identified evidence on the importance of union leaders in shaping unions' effectiveness in achieving their objectives.

The remainder of this paper is organized as follows. Section 2.2 provides institutional background on the CTUM, its role in setting Myanmar's minimum wage, and its member unions. Section 2.3 presents our research design. Section 2.4 presents descriptive evidence on leaders' characteristics and how they compare to their followers in our setting. Section 2.5 discusses the design and results from the consensus-building experiment. Section 2.6 presents the design and results from the mobilization experiment. Finally, Section 2.8 concludes.

2.2 Context

2.2.1 Unions in Myanmar

Unions have been legally allowed in Myanmar since 2011 when the country embarked upon a period of trade liberalization and domestic policy reforms (The Labor Organization Law, 2011). Between 2011-2020, the number of unions grew rapidly. According to the Ministry of Labor, Immigration and Population (MoLIP), as of mid-2020, there were 2,861 registered trade unions.[5] We study unions in Myanmar's export-oriented garment sector, which is the largest exporting industrial sector in Myanmar; as of 2020, approximately 600 factories employed nearly 500,000 workers (Myanmar Garment Manufacturers Association, 2020).

Unions in the garment sector are organized at the factory level and negotiate with management about a number of issues.[6] In a survey of garment factories conducted for Tanaka (2020), the most common topics of negotiation, in order, were pay, working conditions, leave, and working hours. This, in part, motivates our focus on the national minimum wage, as it serves as the reference point in relation to which unions negotiate with management about other components of pay, such as performance incentives. While suggestive, in previous work, we combined Tanaka (2020)'s survey data with administrative data on industrial disputes and documented that garment factories with democratically-selected worker representatives were less likely to experience industrial disputes (Lin et al., 2019). We interpret this as motivating evidence that elected worker leaders may contribute to healthier industrial relations in our setting.

[5]These consist of 2,683 basic organizations, 147 township organizations, 22 state/regional organizations, 8 federations, and 1 confederation. Source here (last accessed on June 22, 2020).

[6]Typically, union membership fees are around 2000 kyats (USD 1.4) per month.

Unions' organizational structure and leaders

The Myanmar Labor Organization Law (2011) sets the terms required to establish a union officially recognized by the law. According to the Law, any group of 30 or more workers can form a factory-level union. To form a union, members must elect seven leaders who form the union's Executive Team (ET). The president leads the union's ET, which also includes an Executive Committee comprising one secretary, one treasurer, and four other elected members (see Figure 2.1). The ET members' tasks differ depending on their position, but one of the key tasks is to regularly attend meetings with the factory management. The basic requirements to become a member of the ET are that a worker has worked at the factory for at least six months, is at least 21 years old, and has a valid national identification number. The Law prescribes that elections are held every two years (unless the president resigns, in which case an emergency election is held). There is no term limit on ET members.

Below the ET, line/team leaders (LLs), play a critical role in facilitating communication with workers.[7] LLs are not elected by union members but are instead recommended by union members, selected by the ET, or self-nominated. Their tasks mostly revolve around communicating with union members as well as recruiting new members.

In our setting, being a union leader is not a paid job. Union leaders are workers in the factories, and evidence from our survey suggests that there are non-trivial costs of becoming a union leader. 70% of presidents and 40% of LLs reported having experienced disadvantages at their factory related to their union activity. Moreover, although the estimates are noisy, Presidents (LLs) seem to face a 20% (15%) wage penalty after controlling for skill measures (average sewing efficiency, number of operations, skill grade), demographics (age, gender, migrant status, education, factory tenure, experience), and personality traits (extraversion, agreeableness, conscientiousness, neuroticism, and openness (Rammstedt and John, 2007)).

The Confederation of Trade Unions in Myanmar (CTUM)

The CTUM is the largest confederation of trade unions in Myanmar. In 2015, the CTUM was officially recognized as the only national-level trade union confederation in Myanmar, marking a significant phase in Myanmar's labor movement. As of 2019, 42 garment factories in Myanmar had a factory-level basic union affiliated to the CTUM, representing 10% of the garment sector and 58% of unions in the industrial sector affiliated to the CTUM.

[7] The CTUM aims to have 1 LL for every 10 workers in unionized factories. In practice, the ratio is smaller; in our sample, the average LL coordinates 33 workers.

2.2.2 The minimum wage in Myanmar

In conformity with the rapidly changing situation of the economy, Myanmar's statutory minimum wage was scheduled to be reconsidered every two years, according to the Minimum Wage Law (2013). A tripartite National Minimum Wage Committee (NMWC) consisting of representatives from employers' and workers' organizations and from the government was responsible for revising the minimum wage. One of the CTUM's key roles was to represent workers in the NMWC. In the 2018 minimum wage negotiations, for example, the CTUM advocated for a 6600 Myanmar Kyat (MMK) (USD 4.87) minimum wage for an eight-hour workday and mobilized workers to demonstrate in favor of its position. The minimum wage was ultimately increased from MMK 3600 (USD 2.65) to MMK 4800 (USD 3.54) and the next revision was scheduled for May 2020.[8]

The minimum wage is highly relevant for garment workers. Appendix Figure 2.6 shows that 59% of workers in our sample reported the legal minimum of MMK 4800 as their daily base wage (not including skill premiums, bonuses, or overtime earnings). Nearly all other workers reported that their daily base wage was just above this amount, and only 4% reported a base wage below it. Turning to daily take-home pay for an 8-hour workday (including base pay, skill premiums, and bonuses), the Figure also shows a dramatic jump up at the legal minimum and that 20% of our sample reported earning between 100-110% of it. In sum, the minimum wage appeared to bind for 20% of our sample, and given its importance in determining base pay, it is plausible that it spilled over to workers earning above this amount (e.g., Autor et al.; Derenoncourt et al. (2016; 2021)).

Raising the minimum wage plausibly entails trade-offs for garment workers. We did not collect data on employers' ability to terminate workers. We do, however, have access to administrative data from the MoLIP on all industrial dispute cases negotiated at the Township Conciliation Body during 2016 in the Yangon region. Out of 407 cases in the garment sector, termination is by far the leading type of dispute (nearly 60% of disputes), followed by wages (nearly 20%). We interpret this as supporting evidence that employers can and do terminate workers. This suggests that, in principle, an increase in the minimum wage could put workers in our sample at risk of job loss.

Against this backdrop, the CTUM aimed to enter the 2020 negotiations equipped with evidence on workers' skills, living costs, and views on the national minimum wage.[9] In 2019,

[8] Due to the seizure of power by the military on February 1, 2021 (the day before the Parliament of Myanmar was due to swear in the members elected at the November 2020 general election), as of late 2021, Myanmar's future democratic prospects are highly uncertain. Anecdotally, union leaders in garment factories have played an active role in mobilizing workers for resistance to the military takeover (source here, last accessed on April 29, 2021).

[9] Due to COVID-19 and the November 2020 elections, there were delays in the minimum wage negotiations and the minimum wage was not revised in 2020.

it sought a collaboration with our research team to collect this evidence, and we agreed to conduct surveys and discussion groups with garment workers. We subsequently produced a joint report with the CTUM on these topics to inform its position. While conducting the surveys and discussion groups, we agreed to run field experiments to understand the role of leaders in shaping collective outcomes.

2.3 Research design

2.3.1 Sampling

We implemented the field activities with workers employed at garment factories in the Yangon and Bago regions that had a factory-level basic union affiliated with the CTUM from December 2019 to March 2020. We focused on these regions because they are home to the majority of garment factories in Myanmar. At the time, 41 garment factories had a union affiliated with the CTUM in these regions. We planned for around 30-35 unions to participate, and our final list included 28 unions.[10] Unfortunately, due to COVID-19, we had to stop our data collection activities, which we explain below, early; 17 unions fully completed the data collection activities, and 19 unions partially completed them.

Appendix Table 2.6 reports summary statistics for the characteristics of the factories in our sample. The average factory size is 1187 workers, and the average union membership rate is 40% of workers. The average number of months the union has been in place at the factory is 29 months, and union presidents' average tenure in the position is 18 months.

Within each factory, we used a sampling protocol that we designed to obtain a sample that was representative of the populations of interest: union leaders (presidents and LLs) and sewing operators (union members and non-members). We focused on sewing operators for two reasons: First, sewing operator is the most common and arguably the most important position type for garment production; sewing operators comprise about 60% of the sewing sector's workforce. Second, the CTUM aimed to collect data on workers' skills, which we supported by developing a skill assessment module for sewing operators. Our assessment method relied on a global industrial engineering database that lists each sewing operation required to produce one piece of a given garment and its complexity in terms of the amount of time required to complete the operation. This type of database does not exist for other parts of the garment production process.

[10] The selection of the unions, and their corresponding factories, was done together with the CTUM. The main criteria were strength of the affiliation to the CTUM, location of the factory with respect to the survey location, and time of union foundation at the factory (some factories were just in the process of finalizing the establishment of the union in the factory).

We conducted a stratified random selection of around 90 workers per factory; within factory, we stratified by line, union membership, and skill level. We detail our random sampling procedure in Appendix 2.14.1. As we discuss below, for each factory, we started the data collection with union leaders and then continued on to the workers. In total, we invited 18 presidents and 1 secretary (19 factories),[11] all of whom participated. We invited 190 LLs from 19 factories, and 170 participated. For workers, due to COVID-19, we only covered 17 factories. We invited 1511 workers and 916 participated (61 % take-up). Among them, we invited 936 union members and 594 participated (63% take-up), and we invited 575 non-union members and 322 participated (56% take-up). Throughout the empirical analysis, we weight observations so that they are representative at the factory level by using probability weights calculated as the total number of workers across factories divided by the number of workers in the specific factory.

Table 2.1 reports summary statistics for union presidents, line leaders, union members, and non-union members. We present the means and standard deviations of various worker characteristics.

2.3.2 Field activities

We embedded a series of experiments in the survey and discussion process. We preregistered the experiments on the AEA's RCT registry. For each factory, we scheduled two consecutive sessions on Sundays. In each session, we included two factories. The sessions were held on Sundays because it is the weekday when most workers have a weekend day. It is important to underscore that participation in the session is costly to workers, as they work very long hours, they only have one weekend day, and for many, they use it to earn an extra wage through overtime work. We compensated participants for their transportation costs (5000 kyats) and time at the average wage rate of a typical working day (6000 kyats).[12] Throughout the activities, we only allowed the research staff and the participants to be onsite when the sessions were taking place; in this way, we aimed to limit any actual or perceived influence of the CTUM on participants' behavior and survey responses.

Figure 2.2 provides an overview of the field activities. In session 1, only presidents and LLs participated. We implemented a survey and a skill assessment as well as a mobilization experiment (EXP 1). The survey covered basic demographic questions as well as information on wages, behavioral characteristics, and psychological traits. The mobilization experiment was about presidents motivating LLs to mobilize workers to attend the session the next Sunday (session 2) and encouraging LLs to produce posters for CTUM's annual International Women's

[11] One union was replacing its president, and the Secretary stepped in the role ad interim.
[12] Many unions preferred to organize communal transportation, in which case we did not reimburse participants.

Day activities (March 8, 2020). Given the much more limited number of presidents compared to LLs, the more limited number of LLs compared to workers, and, crucially, the smaller sample sizes than initially planned due to the Covid-19 outbreak, our results for this experiment are underpowered compared to those with workers. As such, we present this experiment in the Supplementary Materials.

In session 2, which is the main focus of this paper, only LLs and workers participated. In the morning, we implemented a survey, a skill assessment, the public good experiment (EXP 2), and the consensus-building experiment on the minimum wage (EXP 3). The public good experiment was designed to test leaders' potential role of leading by example in the provision of a public good (e.g., Jack and Recalde (2015)). The consensus-building experiment was designed to test how leaders' participation in group discussions about workers' preferred and expected minimum wage levels influenced the group's consensus around these levels.

After lunch, we conducted the mobilization experiment (EXP 4), in which we invited workers to remain for an additional, unanticipated living cost survey for the rest of the afternoon. The CTUM planned to use the living cost data from the survey to campaign for its preferred minimum wage level. We induced a strategic complementarity in turnout at the discussion group level by donating to a worker skills training center for each full discussion group that attended the survey. In this design, we aimed to mirror the incentives faced by workers when deciding whether to participate in collective actions, such as street demonstrations in support of the CTUM's proposed minimum wage level, while avoiding experimentally mobilizing them to engage in potentially risky actions.

Partway through the field activities, we realized that the public good experiment was not working as planned. The endowment of 1500 kyats that we gave to participants (slightly more than USD 1) was too little: Only 7% of leaders and 18% of workers donated less than the full amount (regardless of treatment arm). The censoring of leaders' contributions in particular severely limits the informativeness of the experiment. As such, we do not discuss the public good experiment further in the paper, and we provide more information on it in the Supplementary Materials.

Finally, we also collected audio and video recordings of some of the main activities. Moreover, the field team completed observation forms while running the different activities. When available, we use the data from these sources in our analyses.

2.4 Who are the union leaders?

We conceive of leadership in the spirit of Hermalin (2012), who considers one of the essences of leadership to be the ability to induce others to follow absent the power to compel or to provide formal contractual incentives. This suggests that leaders may exhibit particular characteristics that enable them to influence followers. In this section, we explore this possibility by characterizing how union leaders' traits compare to those of non-leaders. To our knowledge, we provide the first systematic evidence comparing the characteristics of labor leaders and workers who are not leaders drawn from the same population.

As economic theories of leadership are largely silent on the question of who becomes a leader (Hermalin, 2012), we focus on traits that economists identify as relevant for political selection and that psychologists and organizational sociologists associate with individuals' ability to influence collective outcomes. The literature on political selection identifies politicians' ability and their honesty or prosociality as key traits (Caselli and Morelli, 2004). Following this literature, we measure ability using Raven scores (Bilker et al., 2012) and educational attainment. We measure prosociality using altruism.[13] We test for evidence of positive or negative selection into leadership on these traits.

Turning to the psychological and sociological literatures on leadership, a meta-analysis of psychology research on the Big Five Inventory (BFI) personality traits identifies extroversion as the most consistent and highly correlated personality trait with leadership, followed by neuroticism (negative correlation), conscientiousness, and then openness; only agreeableness was not found to be correlated (Judge et al., 2002). The literature also identifies having a strong locus of control (Howell and Avolio, 1993) and grit (Schimschal and Lomas (2018); Caza and Posner (2019)) as important.[14]

Finally, it emphasizes the importance of individuals' charisma, which is considered as a set of behaviors, for leadership ability (House (1977); House and Howell (1992)). Following Antonakis et al. (2016) and Antonakis et al. (2020), we define charisma as the ability to transmit information in a symbolic, value-based, and emotional manner.[15] Of the BFI traits, to our knowledge, only extroversion has been shown to be positively correlated with charisma (Crant and Bateman, 2000). In this section, we focus on leaders' traits, but we return to the concept of charisma in Section 2.5.

[13] We measure altruism via an incentivized question: the respondent chooses how much to keep for herself or donate to a local orphanage institution, out of an endowment of 1500 kyats.

[14] We measure locus of control using a question from the World Values Survey that asks the respondent to indicate using a 5-point Likert scale how much freedom of choice and control the respondent feels to have over the way her life turns out. We measure grit using several questions developed for this purpose by Duckworth and Quinn (2009).

[15] Hermalin (forthcoming) formalizes the notion of charisma in an economic model of leadership.

2.4.1 Qualities of union leaders and non-leaders

We use the following regression specification to compare the characteristics of leaders and non-leaders:

$$Y_{if} = \alpha_0 + \alpha_1 LineLeader_i + \alpha_2 President_i + \gamma_f + \epsilon_{if} \tag{2.1}$$

where Y_{if} is a characteristic of worker i in factory f. $LineLeader_i$ is an indicator of being a line leader, and $President_i$ is an indicator of being a president. γ_f is a factory fixed effect. Finally, ϵ_{if} is the residual. Due to the limited number of clusters (17 factories), we report p-values calculated using the wild cluster bootstrap-t procedure (Cameron et al., 2008).

Table 2.2 presents the results. Each row reports the result from estimating Equation (2.1) for the characteristic in the row. Appendix Table 2.7 presents the same comparisons for each characteristic, estimated including all other variables in the table except the BFI index as controls.[16] Before discussing the characteristics of interest, we mention some demographic and employment differences. Panel A shows that union leaders are significantly less likely to be female and are significantly older than non-leaders. The gender difference is larger for presidents compared to LLs. There is no difference in migration status. Turning to Panel B, union leaders have, on average, 13 (LLs)-19 (presidents) months longer tenures at their factories and substantially more experience in the garment sector. There is not a significant difference between their wages and those of non-leaders.

Turning to leaders' ability and prosociality, beginning with the former, we do not find evidence of selection on ability for LLs. Presidents, however, have significantly higher Raven Scores and more schooling. Further, Appendix Table 2.7 shows that presidents are still positively selected on ability after controlling for their other characteristics, so it's not the case that ability is simply correlated with other traits that predict being a union president. The fact that presidents are higher ability suggests positive selection on this trait through union elections, which goes against the theoretical prediction that the individuals with the highest opportunity cost do not enter into union leadership positions (Caselli and Morelli, 2004). Turning to prosociality, we find that leaders are significantly more altruistic. This is inconsistent with the possibility that individuals pursue union leadership positions to extract rents through dishonest means, and it may help to explain why the leaders in our setting are willing to bear the private costs of leadership discussed in Section 2.2. In sum, the union leaders in our setting are *positively* selected on ability (at least for presidents) and on prosociality, consistent with recent empirical evidence that democratic election of political leaders generates positive selection on ability

[16] For the BFI index regression, we omit the BFI measures as control variables.

(Dal Bó et al., 2017).

Next, we examine leaders' other personality traits. Starting with the BFI, we find a pattern of differences that is highly consistent with the findings of the psychology literature: Leaders are significantly more extroverted, less neurotic, and more conscientious. Interestingly, LLs, whose primary responsibilities entail communication with workers and recruitment of new union members, are significantly more agreeable, but not presidents. Finally, if anything, leaders are less open compared to non-leaders, especially presidents. Reverse-coding neuroticism and taking the average across index components, we find that leaders score significantly higher than workers. We also find that leaders exhibit significantly greater grit and that presidents exhibit significantly greater locus of control. In sum, we confirm that the leaders in our setting have the personality traits that the psychology literature identifies with individuals in leadership roles.

Overall, our evidence is in line with the view of leadership articulated above, which is that it is a phenomenon that exists independent of office or title and that it entails the ability to induce others to voluntarily follow (Hermalin, 2012). We find that union leaders are positively selected on altruism, and union presidents on ability, and that they possess a psychological ability to influence followers.

2.4.2 Focus on line leaders

In the second half of this paper, we present two experiments on LLs' roles in coordinating workers' views and their collective actions. There are a couple of reasons why we chose to focus on LLs in these experiments. First, there are many more LLs than presidents; in our sample, there are 170 LLs compared to 18 presidents. As a result, it was feasible to conduct larger experiments with LLs that are better powered, which is especially true given the ex-post smaller sample sizes due to the Covid-19 pandemic. Second, there is a degree of specialization in presidents' and LLs' leadership roles, and it is LLs who coordinate and motivate workers toward the union's goals on a day-to-day basis.[17] We note that LLs do not possess all of the traits that presidents do. As the political economy and the psychology literatures identify these traits as important for leadership ability, this would push against our finding that LLs have effects on group outcomes. We explore the importance of LLs' resemblance to presidents in Section 2.5.3 when examining the mechanisms through which leaders build consensus in groups.

As the rest of the empirical analysis focuses on LLs, we denote them as *leaders* in the

[17] Appendix Figure 2.7 shows how Presidents and LLs allocate their time differently among various union-related activities. Relative to presidents, LLs spend significantly less of their time communicating with management and meeting with other presidents and significantly more of their time coordinating members, motivating members, and recruiting new members to the union. Panel D of Table 2.1 and Appendix Figure 2.8 also show that workers seek out LLs for advice and social activities more than they seek out presidents.

remainder of the paper.

2.5 Consensus-building experiment

To succeed, social movements must coordinate their members' views and their collective actions. We begin by examining leaders' role in coordinating views. It was important for the CTUM to achieve consensus among workers on their preferred minimum wage, and possibly, its divergence from more probable minimum wage levels, in order for it to determine a credible public position on the planned 2020 minimum wage adjustment and for it to mobilize workers to turn out in support of its position.

Our reading of the theoretical literature on leadership identified several channels through which leaders may contribute to building this consensus. First, leaders may provide information about the state of the world or payoffs that coordinates workers' views (Hermalin (1998); Caillaud and Tirole (2007); Dewan and Myatt (2008)). Second, leaders may influence workers through their communication skills or charisma (Dewan and Myatt (2008); Hermalin (forthcoming); Antonakis et al. (2020).[18] Third, leaders may have social ties with other workers that enable them to influence them (Bandiera et al., 2009). Finally, workers may perceive that leaders' affiliation with their union or with the CTUM endows them with formal authority to make decisions on workers' behalf (Aghion and Tirole, 1997).

These possibilities motivated us to conduct an experiment in which we randomized the presence of a union leader in a discussion about what minimum wage level workers preferred and expected in Myanmar's planned minimum wage adjustment to investigate whether, and if so, how, union leaders helped to build consensus. In this section, we describe this experiment and report our main findings, which are that union leaders build consensus among workers around their unions' ideal minimum wage levels, rather than influence workers' beliefs about the minimum wage level that would result from the negotiations. We provide evidence on the mechanisms through which leaders build consensus around their unions' preferred minimum wage levels.

2.5.1 Experimental design

This experiment took place in Session 2 after workers completed the baseline survey. We stratified workers by their factory and union membership and randomly assigned them to one of three types of discussion groups. In the first type of group, we randomly assigned a leader

[18] As discussed in Section 2.4, leaders score higher than workers on extroversion, the psychological trait that has been found to be positively correlated with charisma (Crant and Bateman, 2000).

from the same factory to participate in the group's discussion. In the second type, motivated by the possibility that leaders primarily influence workers through their social ties with them, we randomly assigned a leader from a different factory, with whom workers are very unlikely to have social ties, to participate in the group's discussion.[19] The final type of group, with no leader participation, is the control condition.

We report balance tests across the three experimental arms in Appendix Table 2.8. While the treatment and control arms are balanced across nearly all tests, there are a few statistical imbalances. When available, we present treatment effects with controls for the baseline value of the outcome variable. We also present results controlling for covariates selected using the post double selection (PDS) lasso (Belloni et al., 2014), which allows us to test our results' robustness to the possibility that chance imbalances between the treatment and control groups influence our estimates.

When forming the groups, we randomized the size of discussion groups (including the leaders) to be either 5 or 6 persons; hence, the leaders' groups are not necessarily larger. The field team implemented the randomized assignment during the worker survey. At the end of the worker survey, they provided workers and leaders with cards that identified their discussion group number. Thus, workers and leaders arrived in the group discussion room concurrently. We did not provide leaders with any specific identification or instructions to lead the discussion, hence allowing for naturally-occurring behavior.

The field team explained to discussion groups that they would discuss the minimum wage. The team provided a brief background of the minimum wage-setting process and its history in Myanmar. The team then explained that the CTUM would prepare a proposal for the government on the minimum wage increase and that the CTUM wanted to gather workers' expectations and opinions to help determine its proposal. Finally, the team told groups that they would have 30 minutes to discuss and requested participants to turn off their cell phones (barring a specific need to keep them on). They also provided the prompt to discussion groups in writing. The prompt's full text is displayed in Appendix 2.14.2.

The field team provided discussion groups with reporting templates and scrap paper to summarize their groups' opinions, which were placed in the center of the discussion group. At the end of the 30 minutes, groups had 5 minutes to summarize their discussion using the templates. The team informed groups that the discussion summaries would be shared with the CTUM to help it to prepare its minimum wage proposal. At the end of the group discussion session, workers and leaders participated in a follow-up survey about their group's discussion

[19] We provide supporting evidence for the assumption that workers are less likely to have social ties with leaders from other factories when discussing mechanisms in subsection 2.5.3.

and their preferences and beliefs about the minimum wage. We recorded and transcribed the audio from the discussions.[20]

2.5.2 Results

We estimate the effects of leaders' participation on convergence to the union's preferred minimum wage, on convergence to the union's expected minimum wage, and on workers' engagement in the discussion. To measure convergence in preferences (beliefs) to the union's preferred (expected) level, we take the average of the preferred (expected) minimum wage among all union leaders within the factory, including the president, measured during the baseline leader survey. We measure the absolute deviation in each worker's view from this average before and after the group discussion. For the external leader arm, we use the average of the external factory's union leaders. We discuss the robustness to alternative approaches to measuring leaders' views in section 2.5.4.

We test (1) the effect of having a leader and (2) the effect of having a leader from one's own factory versus from a different factory. We estimate:

$$Y_i = \alpha_0 + \alpha_1 Leader_i + \mathbf{X}_i'\beta + \epsilon_i \tag{2.2}$$

$$Y_i = \alpha_0 + \alpha_1 OwnLeader_i + \alpha_2 ExternalLeader_i + \mathbf{X}_i'\beta + \epsilon_i \tag{2.3}$$

where Y_i is the outcome for worker i. $Leader_i$ is an indicator for having a leader participate in your group's discussion. X_i is a vector of strata fixed effects and group size fixed effects. Finally, ϵ_i is the residual. For individual-level regressions, we report standard errors clustered by discussion group. For group-level regressions, we report robust standard errors. In equation 2.3, $OwnLeader_i$ is an indicator for having a leader from your own factory in your group, and $ExternalLeader_i$ is an indicator for having a leader from a different factory in your group. When available, we include a control for the baseline value of the dependent variable. We also present the results using the post double selection (PDS) lasso to select control variables (Belloni et al., 2014). The set of potential controls includes all variables in Appendix Table 2.8, personality traits, and psychological traits.

Table 2.3 presents the results. Panel A presents the main effect of having a leader participate, while Panel B presents the effects separately for internal and external leaders. Columns (1)-(2) show that leaders' participation causes workers' *preferences* for the minimum wage to converge to the union's preferred level. There is a 20% decrease in the average absolute deviation from the

[20]Due to an implementation error in the field, discussions for 35 groups were not recorded; consequently, we have transcripts for 167 out of 202 groups.

union's preferred view ($p<0.05$). Interestingly, this effect is not solely driven by union leaders from workers' own factory, although the effect is qualitatively larger for this group; Panel B shows that leaders from external factories induce convergence to their own union's preferred minimum wage. In our placebo and robustness tests, we show that this effect is not an artifact of how we construct the outcome variable. These results support the hypothesis that social ties may matter, but they are not the only channel through which leaders influence followers. This result is consistent with sociological theories of leadership in social movements, which describe a key role for leaders to build consensus on "the world as it should be" among members (Ganz, 2010).

Turning to *beliefs*, we next test whether leaders play a role in terms of conveying information about the likely outcome of the minimum wage-setting process. Columns (3)-(4) show the effect of the deviation from the union's average expected minimum wage. While negative, suggesting less divergence from the union's expectation, the point estimates are small and not statistically significant. This is also true when we split by own versus external leader.

Appendix Figure 2.7 is helpful to interpret this finding; it ranks presidents' and line leaders' time spent on activities according to how presidents spend their time. Compared to presidents, line leaders spend much less time on tasks that may convey insider information about the minimum wage-setting process, such as meetings with management, meetings with leaders in other unions, and going to court. Consequently, the null result may be due to their more specialized leadership role, which does not lead them to acquire information about the likely outcome of the minimum wage setting process. Appendix Figure 2.9 offers an additional explanation. It plots the coefficient of variation within each factory in baseline preferences and beliefs and shows that workers, compared to leaders, exhibit a significantly larger variation in preferences but not in beliefs. This suggests that beliefs were more aligned to start with and there was less room for a change in views.

2.5.3 Mechanisms

How do leaders coordinate workers' views to build consensus around the union's preferred minimum wage level? In this section, we use our data on the nature of groups' discussions and heterogeneous treatment effect (HTE) analysis to explore the channels suggested by economic theory. As discussed above, we identified several possible channels, including providing information about the state of the world or payoffs (Hermalin (1998); Caillaud and Tirole (2007); Dewan and Myatt (2008)); influencing workers through their communication skills or charisma (Dewan and Myatt (2008); Hermalin (forthcoming); Antonakis et al. (2020)); having social ties with other workers that enable them to influence them (Bandiera et al., 2009); and being perceived by workers to have formal authority to make decisions on workers' behalf due to their affiliation with the CTUM (Aghion and Tirole, 1997). The results in Table 2.3 Panel B do not support the possibility that leaders influence workers primarily through their social ties. For this reason, we do not further discuss this mechanism in this section.

We measure workers' engagement in the discussion in two ways. First, we use several questions about workers' enjoyment of and engagement in the group discussion from the follow-up survey to construct a worker-level summary index of engagement. Second, we use the field team's assessment of how active a group discussion is, which we also summarize using a group-level summary index. See Appendix 2.14.3 for the variables included in each index.

Next, we ask how a leader's presence influences workers' *engagement* in group discussions. Columns (5)-(6) show that leaders' participation positively affects workers' self-reported engagement. On average, workers report 0.14 standard deviation (SD) higher engagement when a leader participates in the discussion ($p<0.01$). Columns (7)-(8) show that the field team also rates groups with leaders 0.13 SD higher in terms of having an active discussion ($p<0.05$). In Panel B, columns (5)-(8), we test whether these effects may be driven by social connections with the leader. The estimated effects for leaders from workers' own factory and from an external factory are similar for both measures, suggesting that social ties are unlikely to explain leaders' effects on workers' engagement.

In interpreting these results, it's helpful to recall that we did not inform workers of the presence of a leader in their group. In the follow-up survey, we asked workers whether a union leader participated in the group discussion. In Appendix Table 2.9, we test whether workers in groups with union leaders were more likely to perceive a union leader's presence. We find that workers with leaders in their group were about three times as likely to report the presence of a leader (41 pp increase on a control mean of 22 pp). 74% of workers in internal leader groups detected a leader in their group, while 44% of workers in external leader groups detected a leader in their group.

To provide greater insight into what aspects of worker engagement leaders are influencing, we divide the engagement index into three sub-indexes. The first includes survey questions that measure enjoyment, interest, and how worthwhile the group discussion was. The second includes survey questions that measure the extent to which the group reached a consensus on question prompts. The third includes survey questions that measure the worker's own participation in the group discussion.

Appendix Table 2.10 presents the results. Columns (1) and (9) show that leaders have small, positive effects on workers' enjoyment (p=0.067) and self-reported participation in the discussion

(p=0.203). The largest effect, by far, is on workers' perception that the group achieved agreement on the ideal and expected minimum wage levels; leaders' participation increases self-reported consensus by 0.3 SD ($p \approx 0.000$) (column 5).[21] These findings – that leaders build consensus around their unions' preferred minimum wage levels and that workers perceive this to be the case – while leaders only modestly increase workers' enjoyment and engagement, suggest that leaders are not simply facilitating the discussion in ways that enable workers to reach a consensus. It indicates that there are likely other channels through which leaders influence workers' preferred minimum wage level; we explore several possibilities in the next subsection.

Information. Do leaders provide information that coordinates workers' views? To provide insight into this possibility, we analyze the group discussion transcripts. The transcripts do not include speakers' identities, but we asked the transcription company to identify whether there was (1) a confirmed leader, which is a group member who self-identified as a union leader[22]; (2) a possible leader, which is a group member who was not a confirmed leader and who led the discussion and/or explained the questions and answers. Out of 47 (58) internal (external) leader groups, only 4 (1) had confirmed leaders. This is striking because especially for the external leader arm, it suggests that leaders are not directly introducing their formal role in the union to yield influence. Among the remaining leader groups, nearly all, 41 internal/56 external, had a possible leader identified. Control groups were also coded in this way, and 24 of 62 groups had possible leaders.[23]

We combine the transcript data with information on the group's first preferred (expected) minimum wage level entered in the group discussion reporting form.[24] Among leader groups, we examine whether the speaker who first mentions the value is coded as a possible or a confirmed leader or as a worker. The share of leaders in groups is 19.4%. If the first person to introduce the wage level is random, then it should be a possible leader 19.4% of the time. We find that leaders mention the preferred minimum wage first in 39.2% of groups and the expected minimum wage first in 38.4% of groups. In both cases, we reject that leaders and workers are equally likely to

[21] Interestingly, workers perceive that the group achieved greater agreement on both the ideal and the expected minimum wage levels.

[22] Specifically, whether a group member introduced themself as a line leader/union leader or indicated to workers that they are a member of the union's leadership team.

[23] As one of our goals was to separately identify workers' speech and union leaders' speech, the transcription company was aware of the groups' treatment arm. Consequently, while the large difference in the presence of possible leaders between treatment and control groups is consistent with leaders behaving differently from workers, we cannot rule out that the transcribers were influenced by the knowledge of the groups' treatment status. For this reason, we do not analyze this variable as an outcome.

[24] In this analysis, we focus on subsamples of discussion groups with leaders that satisfy the following criteria: 1) the group reported a preferred (expected) minimum wage level in the group discussion reporting form, 2) at least one person mentioned a preferred (expected) minimum wage level in the transcript, and 3) a possible or a confirmed leader was identified in the transcript. 74 (86) groups meet these criteria for preferred (expected) minimum wages.

mention the minimum wage levels first (p<0.000). Evidently, leaders play a role in introducing what are either influential or preferred values of the minimum wage.

We can also examine how the presence of a leader affects the total amount of speech and the amount of speech by workers using the number of words per person in the discussion transcripts.[25] If workers speak less when leaders are present, it would be consistent with leaders speaking and providing information to influence views; if workers speak more, it could be consistent with leaders speaking and providing information, but also with leaders inducing more information-sharing by workers. We first check whether the presence of a leader affects the total amount of speech. Column (1) of Appendix Table 2.13 shows that groups with leaders discuss weakly less than groups without leaders; although not statistically significant, their discussions are about 14% shorter. Column (5) shows that workers speak significantly less when a leader is present.[26] As the decrease in average worker speech is relatively larger than the decrease in total speech, the results suggest that leaders speak more than workers. We interpret these results as consistent with a key mechanism through which leaders influence views such as information provision, as opposed to simply moderating a discussion among workers to increase the amount of information that individual workers provide.

Leaders' personal traits/charisma. Second, LLs resemble presidents more than workers do (both union and non-union members). In Figure 2.3, we show the cumulative distributions of the predicted probabilities of workers and LLs being similar to presidents using a probit model with demographic variables, personality metrics, and psychological metrics. The horizontal dotted line at 0.5 indicates that while LLs in the bottom half of the similarity distribution are indistinguishable from workers in terms of their characteristics, LLs in the top half of the distribution are distinct from the workers and more closely resemble presidents. We also asked workers and LLs about their aspirations to become elected union leaders; LLs are 8.4 pp more likely to aspire to become an elected leader compared to workers (13% of workers hold this aspiration; p-value of diff. < 0.05).

In our analysis, we use the predicted similarity to the president reported in Figure 2.3 as a measure of leader quality.[27] LLs' predicted similarity to the president is positively correlated with an index comprised of several baseline survey questions that measure LLs' effort for the

[25]The discussions were capped at 30 minutes plus 5 minutes for writing down answers, so this analysis should be interpreted as the amount of speech holding fixed the maximum discussion duration.

[26]To prevent a mechanical negative relationship between a leader's presence and workers' speech, we control for the fixed effects of the number of workers, subtracting 1 from the total group size for treatment groups with confirmed/possible leaders. The worker speech analysis should be interpreted as suggestive, though, as it's possible that speech is misattributed between leaders and workers when coded by the data collection company.

[27]In our analysis, we construct a binary indicator for whether a LL is above the median in their predicted similarity to the president.

union's activities (coeff. = 0.312; p-value < 0.001). It is also positively correlated with LLs' aspirations to become an elected union leader (coeff. = 0.172; p-value < 0.05). Hence, it is reassuring that our measure of LLs' similarity to the president based on their personal traits strongly correlates with revealed preference measures of LLs' engagement in behaviors and holding of aspirations that arguably position them to advance up the union's hierarchy.

We mobilize the concept of charismatic leadership using our measure of leader quality, which is leaders' similarity to presidents (described in Section 2.4.2). This measure reflects traits that arguably correlate with charisma, such as personality and psychological metrics, but may also reflect broader traits associated with moving up the union hierarchy. As such, we interpret the measure as capturing leaders' quality, broadly defined. We partition leaders into above (high-) and below (low-) median similarity types.

Before examining HTEs by leader similarity, we verify that high- and low-similarity leaders in the same factory have the same information about their union's preferred and expected minimum wage levels.[28] We verify that high- and low-similarity leaders are similar in terms of their social ties with workers.[29] Finally, by focusing on leaders at the same tier of their union's hierarchy, we ensure that high- and low-similarity leaders have the same formal authority in their unions.[30]

The definition of charismatic leadership we adopt emphasizes an individual's ability to transmit information in a symbolic, value-based, and emotional manner (Antonakis et al., 2016) (see Section 2.4), while our measure of leader similarity is based on leaders' traits. Consequently, we first test whether high- and low-similarity leaders behave differently, as measured by the research team. Our identification strategy relies on within-factory variation, but our measure of leader quality considers absolute quality; as such, we compare high- and low-similarity leaders' behavior with and without factory FE. Beginning with the absolute comparison (without factory FE), which is the more relevant validation of our measure, Panel D of Appendix Table 2.15 shows that high similarity leaders are rated significantly higher in leadership behaviors almost across

[28] Appendix Figure 2.10 provides evidence in favor of high- and low-similarity leaders having the same information about their union's preferred and expected minimum wage levels. It plots the average baseline views of high- and low-similarity leaders within the same factory. The points are clustered around the 45-degree line in both subfigures, illustrating that both types hold similar views. High- and low-similarity leaders both also have high rates of engagement and interaction with their union. For example, both high- and low-similarity leaders report attending around 9 meetings in the previous 4 months (p-value of diff=0.818). And while high-similarity leaders report meeting with the union president/secretary/treasurer more often, low-similarity leaders also report meeting with them often (6/5/4 times, respectively).

[29] The correlation between the similarity measure and the number of times leaders report socializing with union members outside of union activities in the previous 4 months is 0.013.

[30] In terms of *perceived* authority by workers, Appendix Table 2.9 shows that workers are more likely to perceive the presence of a low-similarity leader. More specifically, this is true for workers assigned to groups with external leaders, while workers' awareness of high- and low-similarity leaders from their own factory is statistically indistinguishable (column 5). If leaders influence group outcomes solely through their authority due to their affiliation with their union or with the CTUM, then we should expect the same effects for high- and low-similarity leaders from workers' own factory. For leaders from an external factory, we should find stronger effects for low- compared to high-similarity leaders.

the board. Evidently, our quality measure based on individuals' traits is highly predictive of their real-world behaviors. Within factory, high-similarity leaders also outperform low-similarity leaders on all measures, although the differences are slightly smaller, and the two measures lose statistical significance. Consistent with high-similarity leaders taking a more active role in the discussion, they are also more likely to be the first speaker to introduce the preferred minimum wage level that appears in the group discussion form (44% compared to 33%), although not for the expected level (40% compared to 37%).

Having established that high-similarity leaders exhibit greater leadership and more actively introduce information about the minimum wage, we test for a role for leader similarity by comparing the effects of high versus low-similarity leaders on our main outcomes. Panel A of Appendix Table 2.16 presents the results. Across all outcomes, the effects are qualitatively larger for high-similarity leaders compared to low-similarity ones. High-similarity leaders decrease the average absolute deviation from the union's preferred minimum wage level by about 24% compared to about 12% for low-similarity leaders (*p*-value of diff=0.213). The effect on the "active group" index is also driven by groups with high-similarity leaders (*p*-value of diff=0.074). Interestingly, both types of leaders increase workers' self-reported engagement, but Appendix Table 2.10 shows that only high-similarity leaders increase self-reported participation (column (11)). This is consistent with the results on participation measured using the group discussion transcripts, which show that low-similarity leader groups drive the negative effect on the amount of discussion (Table 2.13 column (3)) and that low-similarity leaders crowd out workers' speech significantly more than high-similarity leaders do.

While suggestive, we interpret these findings, in particular that high-similarity leaders appear to induce more convergence to the union's preferred minimum wage level despite crowding out workers' participation in the discussion less, as suggestive evidence that high-similarity leaders deliver "better" outcomes for their unions. Leaders' personal traits matter for their ability to influence outcomes.

Finally, we examine the real leaders' similarity relative to placebo control leaders, whom we define in the same way as the first placebo test above. We use the similarity score to partition the control group into high and low placebo leader similarity. Appendix Table 2.18 presents the results. For minimum wage preferences and beliefs, we first use the baseline construction of the outcome (first column, respectively), then we exclude the individual leader's views from the union views (second column, respectively) and finally, we use the assigned leader's baseline views as the reference view (third column, respectively). Across numerous specifications, our main results continue to hold: High-similarity union leaders are the most effective at inducing convergence to the union's preferred minimum wage and increasing engagement in the discussion. We cannot reject that the effects of low-similarity union leaders and placebo leaders are the same, which is consistent with their being indistinguishable in terms of their similarity to union presidents (Figure 2.3). This is consistent with the finding that leader quality, beyond formal authority, matters.

We conduct a robustness test for our leader similarity measure, which is that we drop one family of variables in the prediction model at a time (i.e., demographics, personality traits, psychological traits, and education/tenure) and re-estimate the results.

Organization of the discussion. The results in Table 2.3 indicate that leaders positively affect workers' self-reported enjoyment and participation, so we begin by examining the possibility that leaders induce convergence by acting as discussion coordinators. First, we examine leaders' effects on each variable in the "active discussion" index, which was measured by the field team. Appendix Table 2.12 reports the results. It shows that groups with leaders are 17.3 pps (66%) more likely to have a member observed summarizing opinions and 18.4 pps (28%) more likely to have one taking notes (columns (5) and (6), respectively). Leader groups also have significantly fewer workers who are distracted during the discussion (column (2)). There is no difference, though, in whether a member is actively facilitating the discussion or asking workers' opinions. These results suggest that leaders are combining and organizing workers' views, as opposed to facilitating to increase the overall amount of discussion.

We can directly test whether a leader's presence influences the amount of discussion using the number of words per person in the discussion transcripts.[31] Column (1) of Appendix Table 2.13 shows that groups with leaders discuss weakly less than groups without leaders; although not statistically significant, their discussions are about 14% shorter. In contrast to the previous results, this effect does vary by whether the leader is from the group's own versus an external factory, as it is entirely driven by leaders from workers' own factory (column 2). We return to this finding below.

Leaders appear to be combining and organizing workers' views, but they are also crowding-out workers' speech. Given these effects, we may wonder how leader groups' responses to the question prompts of the possible benefits, harms, and heterogeneous effects associated with increasing the minimum wage compare to those of non-leader groups. We test for differences in the number of words input for each prompt in the group discussion form and find that leader groups input 21% longer responses (Appendix Table 2.14, column (1)). As control groups, on average, input between 12-14 words per question prompt, we interpret the increased word count

[31] The discussions were capped at 30 minutes plus 5 minutes for writing down answers, so this analysis should be interpreted as the amount of discussion holding fixed the maximum discussion duration.

as more substantive responses.[32] In sum, groups with leaders achieve more substantive responses with less discussion and leave workers more satisfied with their experience.

Formal authority. The leaders whom we study have no formal authority or responsibility in the context of the group discussion experiment, but it's still possible that their affiliation with their union or with the CTUM plays a role. We cannot rule out that perceived formal authority contributes to our results; that said, the results suggest that at a minimum, formal authority is not the only mechanism, and more likely, it is not the key mechanism. First, we find that leaders rarely invoke their formal authority in the group discussions. Second, we find that high-similarity leaders have qualitatively and in some cases statistically significant larger effects compared to low-similarity leaders (Table 2.16). This is despite having the same formal authority in the union and despite workers being significantly more likely to perceive the presence of a low-similarity union leader (Table 2.9 column (4)). Further, thus far, we have pooled union members and non-members in our analyses. We might expect, though, that union leaders' formal authority is stronger for members of their organization.[33] We explore this possibility in Appendix Table 2.17, which presents HTEs by union affiliation. We do not find a consistent pattern of heterogeneity by union affiliation.

2.5.4 Placebo and robustness tests

We conducted several placebo and robustness tests for the main results. We summarize the findings of these tests here and provide a thorough discussion of them in Appendix 2.13. In our first placebo test, we identify "placebo leaders" in control groups who resemble actual leaders. We test whether there is greater convergence to the real leader's view in treatment groups compared to the placebo leader's view in control groups, and we find that there is. We then return to the baseline construction of our outcome variable, but for groups assigned to the external leader arm, we test for convergence to the *internal* union leaders' average preferences (beliefs). We find significantly less convergence to the internal union leaders' preferred minimum wage level in the external leader arm. In our robustness tests, we show that the results are robust to using the median of leader views. We show that leaders have effects even conditional on the predicted leader similarity of the workers in their discussion group and that they hold when controlling for the leader or placebo leader similarity. We also show that the results hold if we do not use probability weights in the regressions. Finally, as leaders are somewhat more likely than

[32] Appendix Figure 2.11 displays the most frequent bi- and tri-gram word combinations in the group discussion form responses.

[33] We acknowledge that is also possible that leaders' charisma is organization-specific.

workers to be men (12.9% compared to 3.3%), which is an observable characteristic and affects the group's gender composition, we show that the results are robust to controlling for groups' gender composition.

2.5.5 Consensus-building experiment: discussion

We have three main findings from the consensus-building experiment. First, the participation of a leader in a group dialogue causally affects consensus-building. In particular, it coordinates workers' minimum wage preferences around the union's ideal level. In contrast, we find no effects on convergence in beliefs, which may be because there was already alignment in beliefs at baseline or because the leaders do not appear to have insider information about the potential outcomes of the policy-setting process. Second, consistent with leaders' building consensus, their participation increases both self-reported and observed measures of engagement. Third, leaders' influence is not solely due to their social ties or to their formal authority; leaders introduce information to the discussion, and their personal traits or charisma matter. Further, leaders' influence extends beyond their organization's boundaries, as their participation also affects non-union members.

2.6 Mobilization experiment

Once a social movement achieves consensus around common objectives and tactics, it must coordinate its members' collective actions to accomplish its goals. In the CTUM's case, once it built consensus around an ideal minimum wage level, it needed to coordinate workers to turn out in support of its position. We next examine leaders' role in this process.

Our reading of the theoretical literature on leadership identified three channels through which leaders may mobilize workers. First, leaders may motivate them. In our setting, the national minimum wage policy-setting process would result in uncertain public benefits, but it was common to workers' shared experience that their wages crucially affected their livelihoods. As such, a key role for leaders may be to emotionally appeal to workers to exert effort to influence this process (Ganz, 2010). Second, workers' decision to participate in collective action around the minimum wage has the features of coordination games among individuals with incomplete information, which often have multiple equilibria. In this sense, a key role for leaders may be to select and to communicate the equilibrium to be played (Dewan and Myatt (2008); Akerlof and Holden (2016)). Finally, workers may be aware that leaders will know whether they turn out, which may influence them through two distinct channels. First, leaders may act as enforcers, monitoring workers' behavior and determining and enforcing sanctions on free-riders (Hermalin, 2012). Leaders may also reward good behavior, so workers who aim to increase their involvement or to pursue leadership positions in the union may want to signal their type to the leader (e.g., Ganz (2010)).

We designed an experiment in which leaders attempt to increase individual workers' participation in collective action through these different channels. In this section, we describe this experiment and report our main findings. We find that leaders do *not* simply motivate workers to participate. Instead, leaders appear to play an important role in coordinating workers' actions on an equilibrium that provides higher turnout/participation. Their observation of workers' decisions to participate also appears to increase participation; the evidence favors a signaling channel, in which workers who aim to positively signal their type turn out when they know that the leader will observe their decision.

2.6.1 Experimental design

We aimed to design the experiment to test the channels through which leaders may influence workers' willingness to participate in a high-stakes, real-world collective action to influence the choice of the minimum wage. We faced the empirical challenge, though, that experimentally mobilizing workers to participate in street demonstrations around the minimum wage would subject workers to undue risk. Consequently, we designed the experiment to mirror the incentives that workers face when deciding whether to participate in these types of collective actions while avoiding many of the associated risks.

The experiment entailed three main ingredients. The first ingredient is a costly action: garment workers have a 6-day workweek and often work overtime on the seventh day. Participants in the session had agreed to participate in a half-day session on their one weekend day, but in the experiment, we invited them to participate in an unannounced cost-of-living survey that required them to stay for the rest of the afternoon. Second, there is a common, public good cause, which is the cost of living survey to inform the CTUM's policy position. Third, we create a strategic complementarity in attendance at the group level by announcing that, for each full discussion group that attends the survey, the research team would donate 8000 kyats (about $5.60) to the CTUM Skills Training Centre.[34] We induced this strategic complementarity to increase the incentives for coordination among workers.

The experiment's design follows directly from its theoretical foundations. It entailed a two-level randomization, illustrated in Figure 2.4. First, we stratified discussion groups by factory and consensus-building treatment arm and then randomized them to high or to low mobilization by the leader. In the former, all but one group member were invited by a leader from the group's

[34] The CTUM Skills Training Centre serves all garment workers, not only union members.

factory. In the latter, only one group member was invited by a leader from the group's factory. Within group, we experimentally varied exposure to the three potential leadership channels:

1. **Motivation:** We varied whether workers are invited by a leader versus by the research staff. We provided leaders and research staff with the same invitation script.

2. **Coordination:** We varied whether workers are informed about how many group members are motivated by the leader.

3. **Observation:** We varied whether workers are informed that the leader will observe their decision to participate.

The experiment was implemented as follows: After workers completed the group discussion and follow-up survey, we provided them with lunch. The field team told workers that they would receive their participation payment after lunch, at which time the session would end and a bus would transport workers back to their factory (the meeting point for workers' sharing transportation).

During lunch, the field time prepared the final experiment.[35] At the end of lunch, the field team informed workers that they would be called into a separate room to sign for their payment and provided them with two paper cards: One that included their number in the order in which they would receive the payment, starting from 1 in each discussion group, and one that was a color-code corresponding to their treatment assignment. Workers were not informed about the meaning of the color coding. The field team also requested that workers turn off their cell phones, barring a critical need to keep it on.

In a separate room, the field team informed leaders about the surprise survey session. Among leaders who could stay, the field team randomly assigned two of them to the room where leaders invited workers to stay for the afternoon session, and they were provided with the invitation script. The rest of the leaders were sent to the room where the survey would take place.

After lunch, the field team called workers by their numbers. When workers entered the payment room, they went to the desk corresponding to the color of their card. Each desk was staffed with a member of the field team, and in the leader motivation treatment arms, a leader. The field team member provided the worker with an envelope containing their payment, the worker signed, and the invitation for the afternoon session corresponding to the desk's treatment arm was made. Appendix Section 2.14.4 provides the scripts for each invitation treatment arm.[36]

[35]During lunch, the field team calculated workers' survey incentive payments and implemented the randomized assignment for the mobilization experiment. The field team also randomly assigned the order in which motivated or non-motivated workers would be invited (either all motivated first or all motivated second). Workers ate lunch with their discussion group members in the discussion room.

[36]Note that our implementation ensured that we did not deceive participants.

The research team carefully planned workers' movement from the discussion room to the payment room and then either directly to the afternoon survey room (if they accepted the invitation) or to the bus (if they did not). We also ensured that there were small amounts of buffer time between workers. These aspects of the design were important to prevent information spillovers across workers and were carefully enforced. While they increased the amount of time required to issue the payments, the field team quickly became adept at implementing the procedures. We report the balance table across the experimental arms in Appendix Table 2.19. As our implementation did not involve deception, this resulted in some treatment arms having a lower number of workers as we had to respect the design constraints of the motivation and coordination arms; in particular, in the coordination arm workers were informed about how many workers in their discussion group had been motivated by the leader (i.e. were in the motivation arm).

First, the CTUM's member unions are responsible to turn out workers in support of the CTUM's position, which requires overcoming free-riding problems. As such, a key role for leaders may be Second, For both of these channels, we expect the effects of social pressure to be larger for union members. In the former case, this is because the punishment would be socially enforced, which only works for workers who are closely connected to other group members. In the latter case, this is because we expect that union members care more about how union leaders perceive them compared to non-members.

2.6.2 Results

First, we test for evidence of leaders as motivators and/or as coordinators. If a key role for leaders in our setting is to motivate their followers, then we should find that workers invited to participate in the afternoon session by the leader are more likely to take up the offer. If a key role for leaders in our setting is to coordinate their followers, then we should find that workers informed that they are in a high-coord leader motivation group (i.e., that the leader invited all but one member of their group) should be more likely to take up the offer compared to those who are informed that they are in a low-coord leader motivation group (i.e., that the leader invited one member of their group). We estimate the following model:

$$Y_i = \alpha_0 + \alpha_1 Leader_i + \alpha_2 HighCoord_i + \alpha_3 LowCoord_i + \mathbf{X}_i'\beta + \epsilon_i \tag{2.4}$$

where Y_i is take-up of the afternoon session offer for worker i. $Leader_i$ is an indicator for being motivated by the leader, $HighCoord_i$ is an indicator for being informed that you are in a high coordination group, and $LowCoord_i$ is an indicator for being told that you are in a low coordination group. X_i is a vector of strata fixed effects (factory x discussion group) and treatment assignment for the social pressure arm, which we abstract from for the purpose of presentation. Finally, ϵ_i is the residual. We report robust standard errors. As with the previous experiment, we also present the results using the post double selection (PDS) lasso to select control variables. The set of potential controls includes all variables in Appendix Table 2.19, personality traits, and psychological traits.

It is plausible that leaders influence mobilization through multiple channels and that these channels complement or substitute for each other. We next test whether motivation and coordination by the leader are complements or substitutes. We estimate the following model:

$$Y_i = \alpha_0 + \alpha_1 Leader_i * HighCoord_i + \alpha_2 NoLeader_i * HighCoord_i + \\ \alpha_3 Leader * LowCoord_i + \alpha_4 NoLeader_i * LowCoord_i + \mathbf{X}_i' \beta + \epsilon_i \quad (2.5)$$

where $NoLeader_i$ is an indicator for being invited by the research team (no leader motivation).

Finally, we consider the possibility that the nature of a leader's influence may be general or may be specific to their organization. As such, we test for heterogeneous treatment effects by a worker's union affiliation.

Table 2.4 presents the results. In all columns, the reference group is workers who are invited by the research staff and are not provided with coordination or social pressure information. Columns (1)-(2) show that motivation by the leader does not affect the take-up of the offer; the estimated effect is close to zero and actually slightly negative. Evidently, in this setting, we do not find evidence of a role for motivation through charismatic leadership. That said, we are pooling all leaders; it's possible that our main effects mask heterogeneity by leader type. Unfortunately, we cannot explore this possibility in this experiment, as we do not observe which leader is responsible for inviting a given worker. In any case, unlike the group discussions, the scope for heterogeneity analysis would be limited in this experiment as there were only two LLs per factory inviting workers.

In contrast, high coordination by the leader substantially increases take-up compared to low coordination. Moving from being informed that the leader will invite one group member only to being informed that they will invite all but one group member increases take-up by 13 pp or 38% compared to the control group mean (p=0.084). Evidently, in our setting, leaders do not appear to play a key role as motivators but do appear to play an important coordinating role.

Turning to our test for complementarity or substitution effects, in columns (3)-(4), we see that the effects of moving from low to high coordination by the leader are qualitatively much larger for those who are also motivated by the leader: There is a 26 pp or 75% increase in take-up (p=0.019) compared to a 12 pp or 34% increase in take-up (p=0.267) when not motivated. While motivation by the leader alone may not influence take-up, it does work as a complement to coordination in increasing turnout.

Finally, in columns (5)-(6), we present the results for our test of heterogeneous treatment effects by union affiliation. Beginning with motivation, we find that union members are no more likely to take up the offer when invited by the leader, while non-members are somewhat less likely to take up the offer when invited by the leader. The estimated treatment effects, however, are not statistically significant, nor is the difference in the treatment effect of motivation by the leader between these groups. As such, this evidence should be interpreted as suggestive evidence that motivation by the leader is relatively more important for union members compared to non-members.

Turning to coordination, the effect of moving from low to high coordination is larger and is statistically significant for non-union compared to union members (although the difference between union and non-union members for each respective effect is not statistically significant). While it may be initially puzzling that the response is greater for non-union members, we can explain this finding based on Bayesian updating with normally-distributed priors. In our data, non-union members have lower average priors about their group members' likelihood of participation compared to union members, but with higher variance and a slightly fatter right tail. As such, in the high coordination arm, we expect them to update more positively about the likelihood of their group members' take-up, which could generate the more positive effect. In the low coordination arm, the more negative effect could be driven by the non-union members in the right tail of the distribution of priors. Hence, the potential for leaders to influence coordination appears to be greater for non-union compared to union members due to the underlying differences in priors between the two groups. In short, we conclude that leaders' coordinating role matters more when there is a greater need for coordination (i.e. for non-union members). This evidence is also consistent with the results from the consensus-building experiment, in which we do not find strong evidence of organization-specific charisma.

We next analyze how being informed that a leader will observe their decision affects workers' take-up of the invitation. We estimate the following model:

$$Y_i = \alpha_0 + \alpha_1 SocialPressure_i + \mathbf{X}_i'\beta + \epsilon_i \tag{2.6}$$

where $SocialPressure_i$ is an indicator for being in the social pressure treatment arm. Now, $\mathbf{X}_i'$ is a vector of strata fixed effects and treatment assignments for the motivation and coordination arms.

As discussed above, we identify two potential mechanisms through which observation of the workers' decision by the leader may influence take-up: Leaders acting as judges, sanctioning workers who do not turn out, or workers perceiving that turning out sends a positive signal about their type to the leader. Depending on workers' priors, these mechanisms generate different effects. Under the sanctioning hypothesis, workers with higher priors about their group members' likelihood of accepting the offer should be more likely to take it up when their decision is observed by the leader. In equation 2.7 below, $\alpha_2 < \alpha_1$. Under the signaling hypothesis, workers with lower priors should be more likely to take it up; in this case, $\alpha_2 > \alpha_1$.

$$Y_i = \alpha_0 + \alpha_1 SocialPressure_i \times HighPrior_i + \alpha_2 SocialPressure_i \times LowPrior_i + \alpha_3 HighPrior_i + \mathbf{X}_i'\beta + \epsilon_i \quad (2.7)$$

As we did not directly measure workers' priors, we use a random forest algorithm to predict them using the control group's characteristics and take-up. We implement the random forest algorithm using the randomForest package in R, which is widely used and implements a standard algorithm. We include variables that measure demographics, personality, sociability, employment characteristics, union membership, and group discussion treatment status and engagement. We use the control group as the training set and grow a forest with 250,000 trees; we use the default settings for other parameters, such as the number of variables to randomly sample at each split for growing trees. We stratify the random sampling of control workers by factory. Once we have created the random forest model, we apply it to the rest of the sample to generate each worker's predicted likelihood of take-up. We use these predicted likelihoods to construct, for each worker, the expected probability that all other workers in their group will take up the offer. We then partition the sample at the median into high- and low-predicted priors.

Table 2.5 presents the results. In columns (1)-(2), we show that informing workers that the leader will observe their decision increases take-up by 4.7 pp or about 14% (not statistically significant). In columns (3)-(4), we test whether the effect is heterogeneous by union membership. As discussed above, we hypothesize that the treatment effects should be larger for union members under both potential channels. Indeed, as shown in column (4), the effect is entirely driven by union members, for whom the effect is a 7.0 pp or 21% increase in take-up, while for non-union members, it is small and actually negative. Due to power limitations, we are unable to reject, however, that the effects are the same (p=0.200).

Turning to the potential roles of sanctioning versus signaling, in columns (5)-(6), we present results for workers with high and low priors, respectively. Among workers with above-median priors about their groupmates' likelihood of take-up, there is no effect, which indicates that sanctioning does not appear to be the key channel. In contrast, among workers with below median priors, being told that a leader will observe their decision increases take-up by 10 pp or 30%. Evidently, there is strong evidence in favor of a signaling mechanism in which workers aim to signal their type to the leader in order to increase their prestige or status with the leader.

Finally, in columns (7)-(8), we further explore the social signaling mechanism. Based on the theory, we expect that the effect of being observed by the leader is strongest for union members with low priors. In these columns, we interact the social pressure treatment with having a high or low prior, respectively, and then with an indicator for union membership. We see that the effect is entirely driven by union members with low priors (+15 pp), while there is no effect on non-union members with low priors (p-val of difference <0.0001). Similarly, there is no effect on union members with high priors, and the effect on non-union members with high priors is actually negative.

2.6.3 Mobilization experiment: discussion

We have four main findings from the mobilization experiment. First, leaders' role in mobilizing workers is *not* simply to motivate them to participate. Second, leaders *do* play a key role in coordinating workers to achieve an equilibrium that provides higher turnout/participation. Third, leaders also appear to influence participation through exerting social pressure; in our setting, this influence is limited to members of the leaders' organization (i.e., union members). In principle, this social pressure could take the form of sanctioning bad behavior or rewarding good behavior. Our fourth finding is that we find no evidence of the former, but we do find evidence of the latter: Workers with low priors about their group-mates' take-up are *more* likely to take up the invitation when told that the leader will observe their decision.

2.7 Consensus-building & mobilization

Finally, in Appendix Table 2.11, we test how exposure to a leader influences workers' decisions to take up our surprise invitation for the afternoon session. This surprise invitation happens later in the day, after lunch, during the mobilization experiment. In the mobilization experiment, there is a control group that is *not* mobilized by a leader to take up our invitation. We use workers in this group for this analysis. While our small sample size limits our statistical power, we see that workers who are exposed to a leader during the group discussion are 11 pps more likely to accept our invitation. The effect is large compared to the control group mean, a 33% increase. Again, the effect is similar for leaders from one's own versus an external factory.

2.8 Conclusion

Social movements are critical drivers of institutional change, but to succeed, they must overcome severe collective action problems. Unlike other organizations, however, social movements cannot rely on formal hierarchies and contracts to align incentives and to mobilize members. In the absence of these organizational tools, we identify leaders as playing potentially important roles. We define leaders as individuals who have "...the ability to induce others to follow absent the power to compel or to provide formal contractual incentives" (Hermalin, 2012). While a large theoretical literature has formalized several channels through which leaders may influence collective action, empirical tests in real-world social movements remain scarce.

In this paper, we present experimental evidence on leaders' role in building consensus around common objectives and on mobilizing members to take privately costly actions that convey uncertain public benefits. We study union leaders in Myanmar's garment sector. We first document that union leaders are distinct from union members and non-members along key traits that economists identify as relevant for political selection (Caselli and Morelli, 2004) and that organizational sociologists and psychologists associate with the ability to influence collective outcomes (Judge et al., 2002), respectively. We find that, relative to other workers, leaders are higher ability and substantially more altruistic, are more extroverted, less neurotic, and more conscientious, and they have greater grit and locus of control. They also have more work experience.

Next, we present experimental evidence that leaders shape consensus about the labor movement's objectives by improving group engagement and increasing workers' consensus around their unions' preferred minimum wage levels. We do not find evidence that leaders provide insider information about the potential outcome of the minimum wage policy process. In terms of channels, a leader's personality as opposed to formal authority or social connections is a key driver behind the outcomes.

Finally, we investigate how leaders mobilize workers to undertake a privately costly action for the common good. We find that leaders in our setting play an important role in coordinating workers' participation; in contrast, motivation by leaders alone does not increase participation. Finally, monitoring by leaders also increases take-up through a signaling channel.

2.9 References

Aghion, Philippe and Jean Tirole, "Formal and Real Authority in Organizations," *Journal of Political Economy*, 1997, *105* (1), 1–29.

Akerlof, Robert and Richard Holden, "Movers and shakers," *The Quarterly Journal of Economics*, 2016, *131* (4), 1849–1874.

_ , Anik Ashraf, Rocco Macchiavello, and Atonu Rabbani, "Layoffs and productivity at a Bangladeshi sweater factory," 2020.

Antonakis, John, Giovanna d'Adda, Roberto Weber, and Christian Zehnder, " "Just words? Just speeches?" On the economic value of charismatic leadership," Technical Report 2020.

_ , Nicolas Bastardoz, Philippe Jacquart, and Boas Shamir, "Leadership by example in the weak-link game," *Annual Review of Organizational Psychology and Organizational Behavior*, 2016, *3*, 293–319.

Autor, David H., Alan Manning, and Christopher L. Smith, "The Contribution of the Minimum Wage to US Wage Inequality over Three Decades: A Reassessment," *American Economic Journal: Applied Economics*, 2016, *8* (1), 58–99.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul, "Social connections and incentives in the workplace: Evidence from personnel records," *Econometrica*, 2009, *77*, 1047–1094.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen, "Inference on Treatment Effects after Selection among High-Dimensional Contr," *Review of Economic Studies*, 2014, *81* (2), 608–650.

Bilker, Warren B, John A Hansen, Colleen M Brensinger, Jan Richard, Raquel E Gur, and Ruben C Gur, "Development of abbreviated nine-item forms of the Raven's standard progressive matrices test," *Assessment*, 2012, *19* (3), 354–369.

Bó, Ernesto Dal, Frederico Finan, Olle Folke, Torsten Persson, and Johanna Rickne, "Who becomes a Politician?," *Quarterly Journal of Economics*, 2017, *132* (4).

Bolton, Patrick, Markus K. Brunnermeier, and Laura Veldkamp, "Leadership, Coordination, and Corporate Culture," *Review of Economic Studies*, 2012, (0), 1–26.

Boudreau, Laura, "Multinational enforcement of labor law: Experimental evidence from Bangladesh's apparel sector," *Working Paper*, 2021.

Breza, Emily, Supreet Kaur, and Nandita Krishnaswamy, "Scabs: The social suppression of labor supply," Technical Report, National Bureau of Economic Research 2019.

Cagé, Julia, Anna Dagorret, Pauline Grosjean, and Saumitra Jha, "Heroes and Villains: The Effects of Combat Heroism on Autocratic Values and Nazi Collaboration in France," *Working Paper*, 2020.

Caillaud, Bernard and Jean Tirole, "Consensus building: How to persuade a group," *American Economic Review*, 2007, *97* (5), 1877–1900.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-based improvements for inference with clustered errors," *The Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Cantoni, Davide, David Y Yang, Noam Yuchtman, and Y Jane Zhang, "Protests as strategic games: Experimental evidence from Hong Kong's antiauthoritarian movement," *Quarterly Journal of Economics*, 2019, *134*, 1–57.

Caselli, Francesco and Massimo Morelli, "Bad politicians," *Journal of Public Economics*, 2004, *88.*

Caza, Arran and Barry Z. Posner, "How and when does grit influence leaders' behavior?," *Leadership and Organization Development Journal*, 2019, *40* (1).

Corradini, Viola, Lorenzo Lagos, and Garima Sharma, "Collective Bargaining for Women: How Unions Can Create Female-Friendly Jobs," Technical Report 2022.

Crant, J. Michael and Thomas S. Bateman, "Charismatic leadership viewed from above: the impact of proactive personality," *Journal of Organizational Behavior*, 2000, *21.*

Derenoncourt, Ellora, Francois Gérard, Lorenzo Lagos, and Claire Montialoux, "Racial Inequality, Minimum Wage Spillovers, and the Informal Sector," 2021.

Dewan, Torun and David P Myatt, "The qualities of leadership: Direction, communication, and obfuscation," *The American Political Science Review*, 2008, *102*, 351–368.

Dippel, Christian and Stephan Heblich, "Leadership and Social Movements: The Forty-Eighters in the Civil War," *American Economic Review*, 2021, *111*, 1–35.

Duckworth, Angela Lee and Patrick D Quinn, "Development and validation of the Short Grit Scale (GRIT–S)," *Journal of personality assessment*, 2009, *91* (2), 166–174.

Englmaier, Florian, Stefan Grimm, Dominik Grothe, David Schindler, and Simeon Schudy, "The value of leadership: Evidence from a large-scale field experiment," Technical Report 2022.

Enikolopov, Ruben, Alexey Makarin, and Maria Petrova, "Social media and protest participation: Evidence from Russia," *Econometrica*, 2020, *88* (4), 1479–1514.

Freeman, Richard B, "Labor regulations, unions, and social protection in developing countries: Market distortions or efficient institutions?," *Handbook of development economics*, 2010, *5*, 4657–4702.

Ganz, Marshall, "Leading Change: Leadership, Organization, and Social Movements," in Nitin Nohria and Rakesh Khurana, eds., *Handbook of Leadership Theory and Practice: A Harvard Business School Centennial Colloquium*, Harvard Business School Publishing Corporation, 2010, chapter 19.

González, Felipe, "Collective action in networks: Evidence from the Chilean student movement," *Journal of Public Economics*, 2020, *188*, 1–13.

Grossman, Guy and Delia Baldassarri, "The impact of elections on cooperation: Evidence from a lab-in-the-field experiment in Uganda," *American journal of political science*, 2012, *56* (4), 964–985.

Hermalin, Benjamin E., "Toward an Economic Theory of Leadership: Leading by Example," *The American Economic Review*, 1998, *88* (5), 1188–1206.

Hermalin, Benjamin E, "Leadership and corporate culture," *Handbook of organizational economics*, 2012, pp. 432–78.

_ , "At the helm, Kirk or Spock? The pros and cons of charismatic leadership," *American Economic Journal: Microeconomics*, forthcoming.

House, Robert J, *A 1976 theory of charismatic leadership.* 1977.

_ **and Jane M. Howell**, "Personality and Charismatic Leadership," *Leadership Quarterly*, 1992, *3* (2).

Howell, Jane M. and Bruce J. Avolio, "Transformational Leadership, Transactional Leadership, Locus of Control, and Support for Innovation: Key Predictors of Consolidated-Business-Unit Performance," *Journal of Applied Psychology*, 1993, *78* (6).

Jack, B Kelsey and María P Recalde, "Leadership and the voluntary provision of public goods: Field evidence from Bolivia," *Journal of Public Economics*, 2015, *122*, 80–93.

Judge, Timothy A., Joyce E. Bono, Remus Ilies, and Megan W. Gerhardt, "Personality and Leadership: A Qualitative and Quantitative Review," *Journal of Applied Psychology*, 2002, *87* (4).

Komai, Mana, Philip J. Grossman, and Trevis Deters, "Leadership and Information in a Single Shot Collective Action Game: An Experimental Study," *Managerial and Decision Economics*, 2010, *32* (2), 119–134.

Lin, Min Zarni, Mari Tanaka, Virginia Minni, Hanh Nguyen, Khin Suu Thet, and Rocco Macchiavello, "Industrial relations and workplace communication in Myanmar garment sector," 2019.

Loeper, Antoine, Jakub Steiner, and Colin Stewart, "Influential opinion leaders," *The Economic Journal*, 2014, *124* (581), 1147–1167.

Macchiavello, Rocco, Andreas Menzel, Atonu Rabbani, and Christopher Woodruff, "Challenges of Change: An Experiment Training Women to Manage in the Bangladeshi Garment Sector," *Working Paper*, 2020.

Manacorda, Marco and Andrea Tesei, "Liberation technology: Mobile phones and political mobilization in Africa," *Econometrica*, 2020, *88* (2), 533–567.

Manski, Charles F, "Identification of endogenous social effects: The reflection problem," *The review of economic studies*, 1993, *60* (3), 531–542.

Porta, Donatella Della and Mario Diani, *Social movements: An introduction (3rd ed.)*, Oxford: Blackwell, 2020.

Potters, Jan, Martin Sefton, and Lise Vesterlund, "Leading-by-example and Signaling in Voluntary Contribution Games: An Experimental Study," *Economic Theory*, 2007, *33* (1), 169–182.

Rammstedt, Beatrice and Oliver P John, "Measuring personality in one minute or less: A 10-item short version of the Big Five Inventory in English and German," *Journal of research in Personality*, 2007, *41* (1), 203–212.

Sahin, Selhan Garip, Catherine Eckel, and Mana Komai, "An Experimental Study of Leadership Institutions in Collective Action Games," *Journal of the Economic Science Association*, 2015, *1* (1), 100–113.

Schimschal, Sarah E. and Tim Lomas, "Gritty Leaders: The Impact of Grit on Positive Leadership Capacity," *Psychological Reports*, 2018, *122* (4).

Tanaka, Mari, "Exporting Sweatshops? Evidence from Myanmar," *Review of Economics and Statistics*, 2020, *102*, 442–456.

Visser, Jelle et al., "Trade unions in the balance," *ILO ACTRAV Working Paper*, 2019.

2.10 Figures

Figure 2.1: Union Organizational Structure

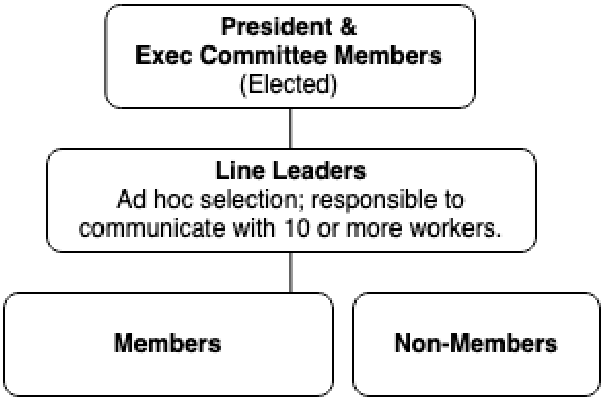

Figure 2.2: Overview of field activities

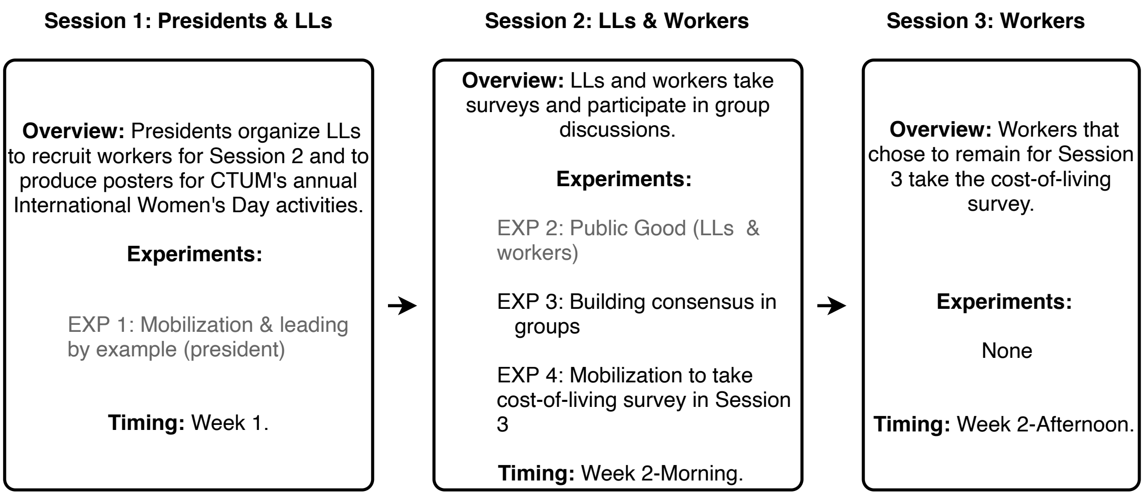

Figure 2.3: Workers' and line leaders' similarity to presidents

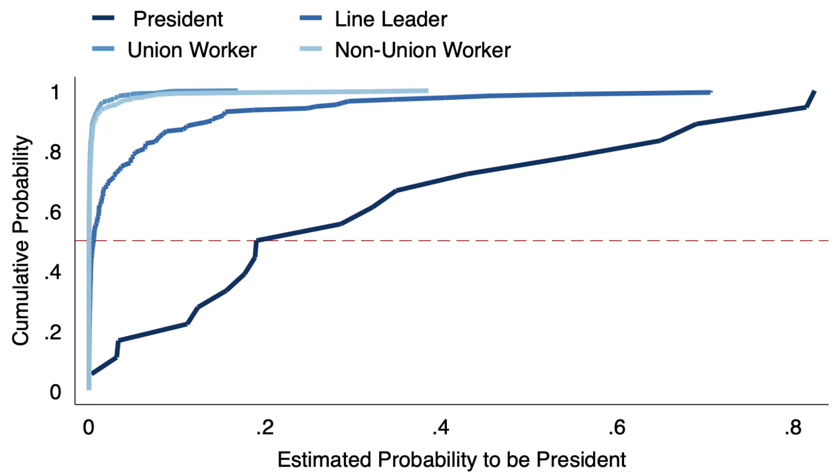

The figure shows the cumulative distribution of the probability of a worker being a president estimated by a probit model with demographic controls (gender, age, education migrant(0/1), months in factory/sector), personality metrics (extraversion, agreeableness, conscientionsness, neuroticism, openness) and psychological metrics (raven score, grit, altruism, choice in life).

Figure 2.4: Mobilization Experiment

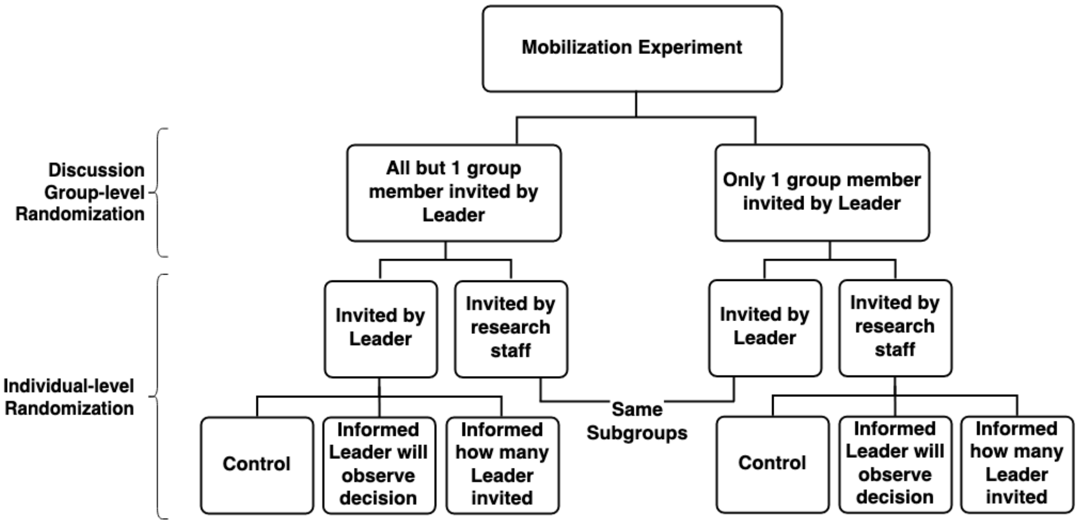

Notes. This figure presents the design of the mobilization experiment. To be completed.

Figure 2.5: Average convergence to union ideal or belief & share mobilized

(a) Union Ideal

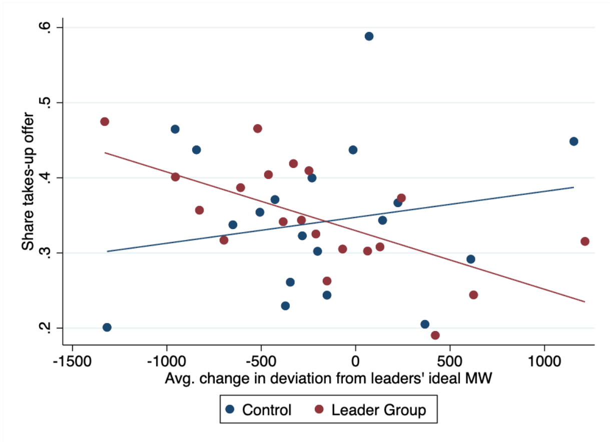

(b) Union Belief

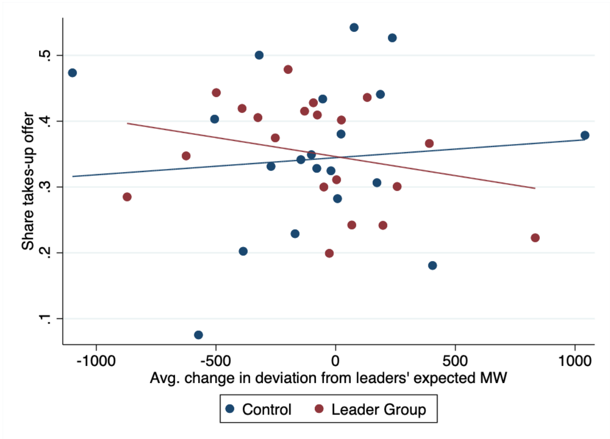

Notes. Both variables have been residualized by factory fixed effects (FEs) and group size FEs, and the mean of each variable has been added back prior to plotting. Group weights are applied.

2.11 Tables

Table 2.1: Summary Statistics

	Presidents	Line Leaders	Union Workers	Non-Union Workers	Total
Panel A: Demographics & Ability					
Female	0.444	0.853	0.963	0.967	0.963
	(0.511)	(0.355)	(0.188)	(0.178)	(0.189)
Age	29.89	26.72	24.75	25.09	24.88
	(6.192)	(6.402)	(5.838)	(6.460)	(6.058)
Migrant	0.444	0.488	0.522	0.518	0.520
	(0.511)	(0.501)	(0.500)	(0.500)	(0.500)
Education (Yrs)	8.500	7.565	7.705	7.854	7.753
	(3.552)	(2.442)	(2.779)	(2.680)	(2.744)
Panel B: Employment & Minimum Wage Views					
Raven Score	6.333	4.429	4.505	4.853	4.620
	(2.275)	(2.581)	(2.743)	(2.844)	(2.778)
Months in Factory	46.50	40.79	33.19	20.60	29.18
	(37.17)	(35.47)	(35.17)	(27.10)	(33.29)
Months in Sector	76.39	73.67	51.63	43.00	49.08
	(63.97)	(57.64)	(48.35)	(50.33)	(49.34)
Income	252464.7	240673.1	238449.6	236588.8	237880.8
	(68261.6)	(38743.7)	(37718.9)	(40663.0)	(38745.6)
Min Wage Ideal	7494.4	7635.3	9338.2	7507.3	8719.6
	(763.5)	(2289.1)	(21030.9)	(2150.6)	(17164.7)
Min Wage Guess	6388.9	6352.7	6484.2	6470.7	6478.3
	(1039.2)	(1123.9)	(1198.1)	(1152.9)	(1181.9)
Panel C: Personality Traits					
Grit	3.754	3.412	2.574	2.558	2.579
	(0.592)	(0.541)	(0.484)	(0.534)	(0.511)
Altruism	1405.6	1406.5	1270.7	1235.9	1261.0
	(279.6)	(292.1)	(418.6)	(483.3)	(439.8)
Locus of Control	4.389	4.200	4.028	4.018	4.027
	(0.698)	(1.200)	(1.265)	(1.376)	(1.301)
Extraversion	3.861	3.644	3.378	3.421	3.396
	(0.763)	(0.815)	(0.754)	(0.716)	(0.743)
Agreeableness	3.972	4.065	3.867	3.900	3.880
	(0.977)	(0.765)	(0.770)	(0.831)	(0.791)
Conscientiousness	4.444	4.168	3.901	4.069	3.959
	(0.511)	(0.715)	(0.769)	(0.778)	(0.775)
Neuroticism	2	2.418	2.647	2.675	2.653
	(0.985)	(0.873)	(0.818)	(0.883)	(0.840)
Openness	2.500	2.932	2.956	3.014	2.974
	(0.686)	(0.757)	(0.754)	(0.783)	(0.764)
Panel D: Communication					
Socialized with union members	6.278	2.029	0.884	0.417	0.750
	(9.228)	(3.237)	(1.655)	(1.095)	(1.583)
Consulted by union workers	4.222	6.947	.	.	6.686
	(3.264)	(16.93)	(.)	(.)	(16.15)
Consulted by non-union workers	2	4.106	.	.	3.904
	(2.449)	(12.84)	(.)	(.)	(12.25)
Observations	18	170	594	322	1104

Notes. Unit of observation is worker. The table summarizes basic demographic characteristics by type of worker. *Education* range from 0 (no education) to 15 (Bachelor's degree). *Income* is the self-reported last month's income in Myanmar kyat. *Socialized with union members* is number of times union leaders and members met other union members for social activities in the past 4 months. *Consulted by union/non-union workers* is number of times union leaders were consulted about issues at the factory in the past month. *Altruism* is amount donated to local orphanage out of an initial endowment of 1500kyats. Probability weights used for the workers columns.

Table 2.2: Differences between Leaders and Workers

	Observations	Worker Mean	Coeff. on Line Leader	Coeff. on President	p-value of diff, cols (3)-(4)
	(1)	(2)	(3)	(4)	(5)
Panel A: Demographics & Ability					
Female	1104	.967	-0.116	-0.518	0.006
			[0.000]	[0.001]	
Age	1104	25.005	1.859	4.918	0.044
			[0.000]	[0.003]	
Migrant	1104	.52	-0.046	-0.085	0.728
			[0.286]	[0.435]	
Education(Yrs)	1104	7.754	-0.176	0.799	0.251
			[0.437]	[0.330]	
Raven Score	1104	4.524	-0.085	1.749	0.003
			[0.734]	[0.005]	
Panel B: Employment & Minimum Wage Views					
Months in Factory	1104	29.888	13.010	18.573	0.327
			[0.000]	[0.000]	
Months in Sector	1104	50.621	24.796	28.216	0.784
			[0.000]	[0.013]	
Income	1103	243154.0	1647.2	13705.9	0.549
			[0.585]	[0.477]	
Preferred Min Wage	1104	7504.258	28.294	171.402	0.566
			[0.823]	[0.482]	
Expected Min Wage	1104	6545.961	-140.598	-91.844	0.841
			[0.085]	[0.692]	
Panel C: Personality Traits					
Altruism	1104	1268.777	142.460	147.861	0.952
			[0.000]	[0.129]	
Extraversion	1104	3.392	0.244	0.488	0.200
			[0.000]	[0.016]	
Agreeableness	1104	3.862	0.214	0.113	0.700
			[0.004]	[0.603]	
Conscientiousness	1104	3.979	0.225	0.507	0.037
			[0.000]	[0.001]	
Neuroticism	1104	2.665	-0.290	-0.670	0.145
			[0.000]	[0.018]	
Openness	1104	3.001	-0.065	-0.473	0.023
			[0.338]	[0.008]	
BFI Index	1104	2.314	0.182	0.261	0.458
			[0.000]	[0.026]	
Grit	1104	2.571	0.854	1.202	0.015
			[0.000]	[0.000]	
Locus of Control	1104	4.008	0.192	0.349	0.435
			[0.081]	[0.079]	

Notes: Unit of observation is worker. Probability weights used. Controlling for Factory FE. p-values calculated using the wild cluster bootstrap-t method.

Table 2.3: Group Discussions: consensus-building

	Deviation from Union Preference		Deviation from Union Belief		Engagement		Active Group	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Leader								
Leader	-223.0**	-223.5**	-19.82	-28.92	0.148***	0.137***	0.131**	0.131**
	(94.63)	(92.95)	(55.04)	(52.62)	(0.0462)	(0.0425)	(0.0547)	(0.0519)
R-squared	0.246	0.227	0.340	0.306	0.090	0.126	0.393	0.394
Panel B: Own versus External LL								
External Leader	-182.9	-185.9*	-11.05	-25.73	0.128*	0.122**	0.139*	0.139*
	(111.2)	(109.1)	(76.99)	(74.61)	(0.0669)	(0.0604)	(0.0760)	(0.0719)
Own Leader	-248.9**	-247.7**	-25.43	-30.94	0.161***	0.147***	0.126**	0.126**
	(105.4)	(103.6)	(64.38)	(61.03)	(0.0483)	(0.0446)	(0.0628)	(0.0594)
R-squared	0.247	0.228	0.340	0.306	0.090	0.126	0.393	0.393
Control Mean	1130.078	1130.078	712.308	712.308	-0.039	-0.039	0.127	0.127
Number of obs.	914	914	914	914	914	914	202	202
p-values								
Own Leader = External Leader	0.532	0.551	0.869	0.951	0.621	0.681	0.877	0.870
PDS lasso selected controls	N	Y	N	Y	N	Y	N	Y

Notes. Unit of observation is worker in all columns but in Col. 7 and Col.8, where it is discussion group.*Engagement* is an index of the following self-reported survey measures of participants' engagement: group consensus on minimum wage prediction/preferences (*GroupAgreePrediction*, *GroupAgreeIdeal*), freedom to express views (*GroupExpressIdeas*, *GroupUnease*[reverse]), interest and enjoyment of the discussion (*GroupInterested*, *GroupEnjoy*, *GroupWaste*[reverse]) and active participation by all members (*GroupAllParticipate*). *Active Group* is an index created from research staff observations which assess group behavior (*ShareEngaged, ShareDistracted, ActiveFacilitation AskingOpinions SummerizingOpinions TakingNotes*). The dependent variables in col. 1, col.2 and col.3, col.4 represent the deviation from the factory average of baseline leaders' preferences and views respectively. Probability weights and standard errors clustered at the group level. Controlling for group size FE and stratification FEs (Factory FEs x Union FEs). Selected controls for Col.6 are *Grit* and *Agreeableness squared (BFI)*. *Deviation from Union Preference* is selected for Col.2. No controls are selected for Col.4 and Col.8. R-squared for columns that applied PDS lasso selected controls are estimated by the correlation between the observed outcome and the predicted outcome.

Table 2.4: Mobilization Session 3, motivation and coordination

	Take-up of surprise offer to participate in survey					
	All		by Leader Invitation		by Union	
High Coord.	0.0790	0.0790				
	(0.0656)	(0.0577)				
Low Coord.	-0.0514	-0.0514				
	(0.0641)	(0.0563)				
Leader	-0.0135	-0.0135	0.0169	0.0157		
	(0.0436)	(0.0384)	(0.0750)	(0.0656)		
High Coord., No Leader			0.101	0.100		
			(0.114)	(0.1000)		
Low Coord., No Leader			-0.0170	-0.0175		
			(0.0778)	(0.0683)		
High Coord., Leader			0.0735	0.0778		
			(0.0777)	(0.0685)		
Low Coord., Leader			-0.178	-0.178*		
			(0.112)	(0.0980)		
High Coord., Union					0.0589	0.0321
					(0.0795)	(0.0688)
Low Coord., Union					-0.0293	-0.0449
					(0.0809)	(0.0705)
High Coord., Non-Union					0.112	0.129
					(0.0956)	(0.0823)
Low Coord., Non-Union					-0.0916	-0.0821
					(0.0954)	(0.0827)
Leader, Union					0.0174	0.0197
					(0.0540)	(0.0471)
Leader, Non-Union					-0.0768	-0.0688
					(0.0675)	(0.0584)
Union					-0.0515	-0.0207
					(0.0750)	(0.0664)
R-squared	0.332	0.311	0.334	0.313	0.336	0.335
Control Mean	0.341	0.341	0.341	0.341	0.341	0.341
Number of obs.	790	790	790	790	790	790
p-values						
Low Coord. = High Coord.	0.130	0.0844				
No Leader, Low Coord. = High Coord.			0.332	0.267		
Leader, Low Coord. = High Coord.			0.0430	0.0194		
Union, Low Coord. = High Coord.					0.395	0.396
Non Union, Low Coord. = High Coord.					0.0880	0.0386
PDS lasso selected controls	N	Y	N	Y	N	Y

Notes. Unit of observation is worker. Probability weights and robust standard errors used. Dependent variable is an indicator for whether worker shows up to take the minimum wage survey. Stratification FEs are included: Factory FEs x Discussion Group FEs.*MDE* is *0.105* for *Invited by Leader*; *0.112* for *Invited by LeaderxHigh Coord.* and *Low Coord.*; *0.169* for *Invited by LeaderxLow Coord.* and *High Coord.*. MDE is determined from power calculations using planned sample size of 1792 workers, 358 discussion groups, 308 LL, and 28 unions, at a 10% significance level and 80% power. No controls are selected for Col.2. *Number of times union leaders and members met other union members for social activities (squared)* is selected for Col.4. *Standardized Months In Factory*, *Standardized Relationship with Managers*,*Raven Score squared*, *Number of times union leaders and members met other union members for social activities (squared)* are selected for Col.6. R-squared for columns that applied PDS lasso selected controls are estimated by the correlation between the observed outcome and the predicted outcome.

Table 2.5: Mobilization Session 3, social pressure

	Base		Cov = Union			Cov = High Prior		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Social Pressure	0.0474	0.0474						
	(0.0454)	(0.0399)						
Social Pressure, Cov=1			0.0766	0.0701	-0.0147	-0.0147		
			(0.0565)	(0.0495)	(0.0392)	(0.0385)		
Social Pressure, Cov=0			-0.0198	-0.0338	0.102***	0.102***		
			(0.0757)	(0.0646)	(0.0382)	(0.0384)		
Social Pressure*High Prior, Cov=1							0.00992	-0.00160
							(0.0481)	(0.0437)
Social Pressure*Low Prior, Cov=1							0.149***	0.147***
							(0.0462)	(0.0467)
Social Pressure*High Prior, Cov=0							-0.0801	-0.0718
							(0.0626)	(0.0599)
Social Pressure*Low Prior, Cov=0							0.00351	-0.0277
							(0.0613)	(0.0580)
R-squared	.33	.31	.33	.33	.34	.32	.35	.35
Control Mean	.34	.34	.34	.34	.34	.34	.34	.34
Number of obs.	790	790	790	790	790	790	790	790
p-values								
Social Pressure (Cov=0) = Social Pressure (Cov=1)			.31	.2	.031	.029		
Union: Social Pressure Low Prior = Social Pressure High Prior							.036	.02
Non Union: Social Pressure Low Prior = Social Pressure High Prior							.32	.58
PDS lasso selected controls	N	Y	N	Y	N	Y	N	Y

Notes. Unit of observation is worker. Probability weights used. Robust standard errors in Columns 1-2 and bootstrap standard errors in Columns 3-4. Dependent variable is an indicator for whether worker shows up to take the minimum wage survey. Stratification FEs are included: Factory FEs x Discussion Group FEs. The *MDE* for Social Pressure is *0.105*. MDE is determined from power calculations using planned sample size of 1792 workers, 358 discussion groups, 308 LL, and 28 unions, at a 10% significance level and 80% power. *PDS* indicates that post-double lasso control selection procedure is applied. *Raven Score squared* is selected for Col.4 and Col.8. No controls are selected for Col.2 and Col.6. R-squared for columns that applied PDS lasso selected controls are estimated by the correlation between the observed outcome and the predicted outcome.

2.12 Appendix: Additional figures and tables

Figure 2.6: Wage distribution

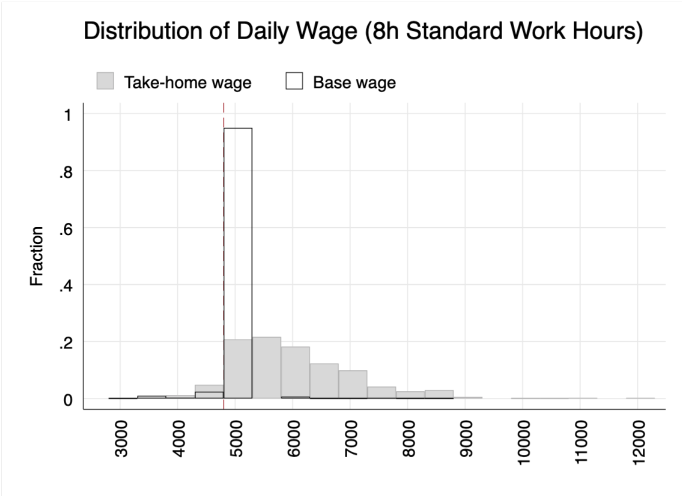

Notes. This figure shows the distributions of daily base wages and daily take-home wages for 8 standard hours for workers in our sample. The transparent bars are the histogram for daily base wages, while the gray bars are for daily take-home wages. The vertical line indicates 4800 kyats, the current minimum wage since 2018. The daily base wage is the base level of the wage for 8 standard hours without reflecting skill premiums, bonuses, and overtime earnings. We calculate the daily take-home wage, which is defined as the daily wage rate for 8 standard hours including the base wage, skill premiums, and bonuses. It does not include overtime work earnings.

Figure 2.7: Time spent on union-related activities

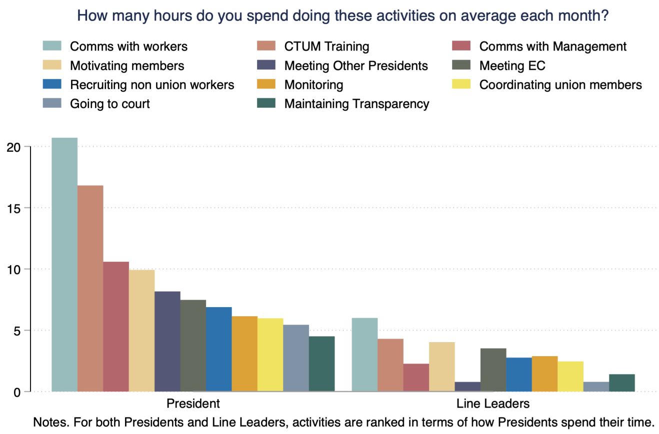

Notes. This figure shows the self-reported time use across different union-related activities separately for presidents and line leaders.

Figure 2.8: Presidents and line leaders' contact with workers

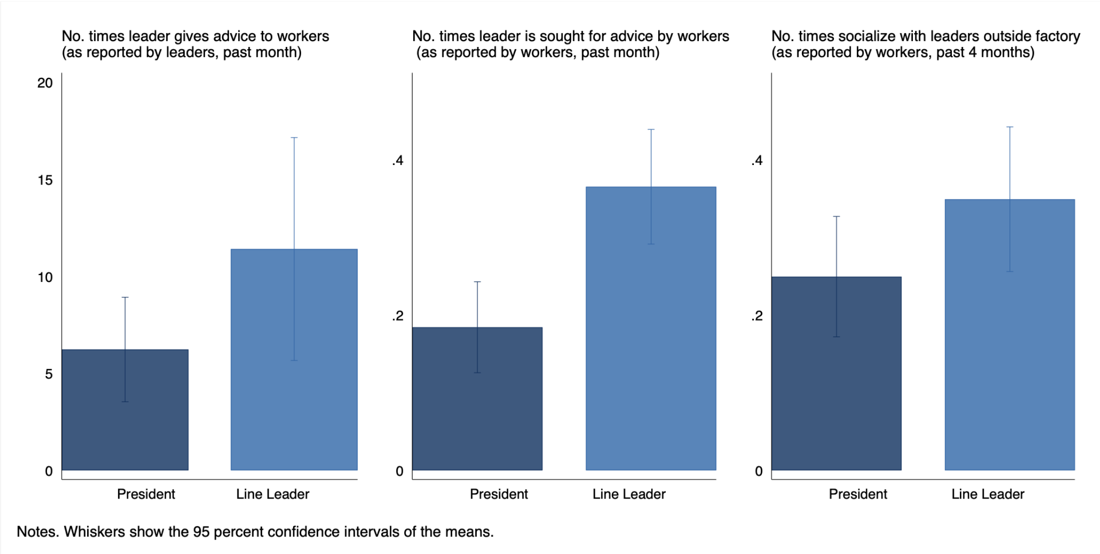

Notes. This figure plots the mean of different self-reported measures of direct contact with workers separately for presidents and line leaders. Whiskers show the 95 percent confidence interval.

Figure 2.9: Coefficient of variation, preferences and beliefs at baseline

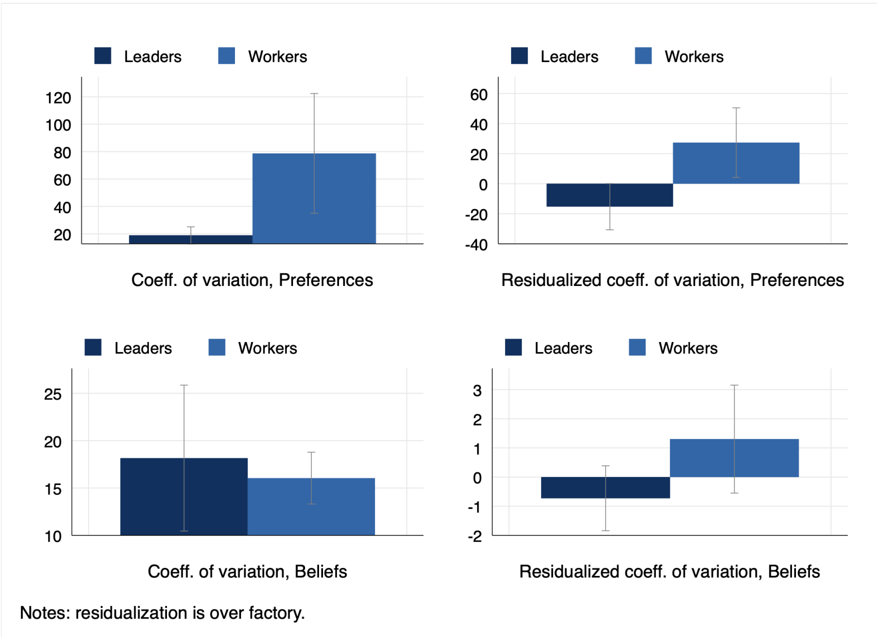

Notes. This figure plots the coefficient of variation in baseline views within factory separately for workers and leaders. Whiskers show the 95 percent confidence intervals, using the variation across factories.

Figure 2.10: High- and low-similarity leaders' minimum wage views at baseline

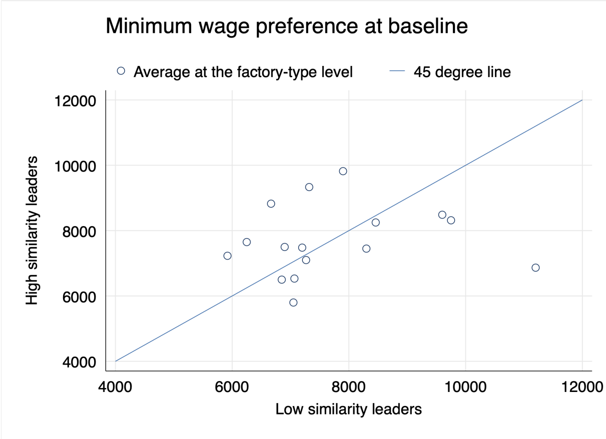

Notes. This figure shows the minimum wage preferences and beliefs at baseline for line leaders, averaged at the factory level. The y-axis is for high-similarity leaders and the x-axis is for low-similarity leaders. The 45-degree line is included.

Figure 2.11: Common bi- and tri-grams in responses to each discussion prompt

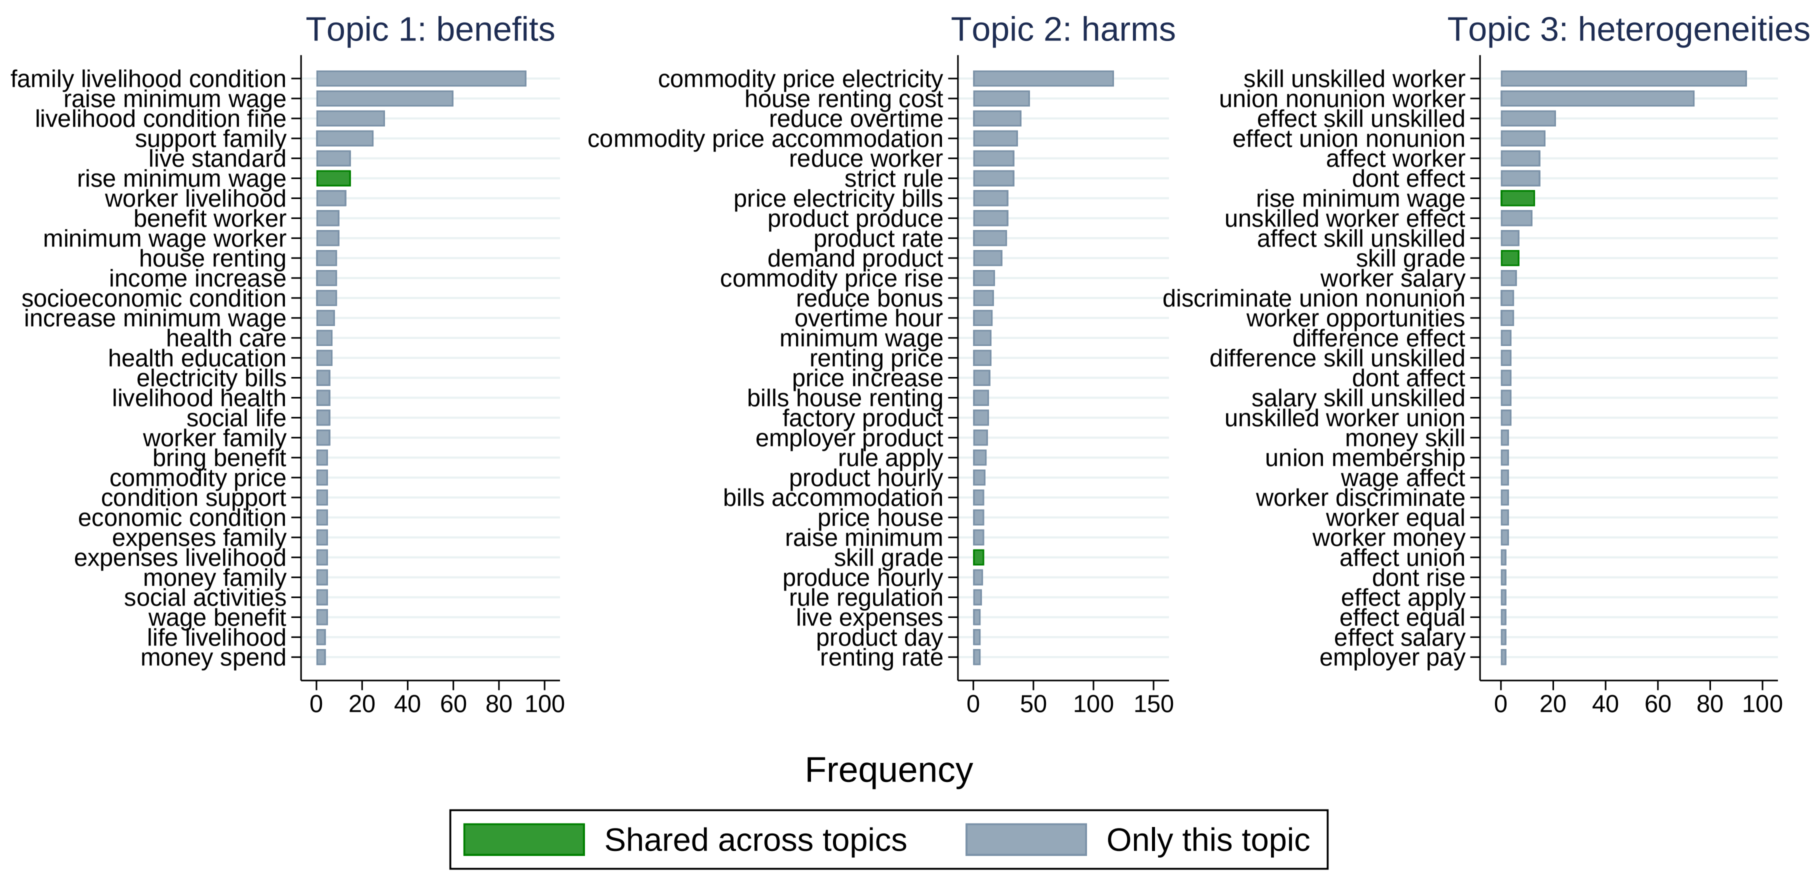

Notes:

Table 2.6: Factory/Union-level descriptive statistics

Variable	**Mean**	**Std. Dev.**	**Min.**	**Max.**	**N**
Number of Workers	1187.5	673.3	450.0	2860.0	17
Number of Union Members	505.8	426.0	100.0	1938.0	17
Proportion Unionized	0.4	0.2	0.1	0.8	16
Female Union President	0.5	0.5	0.0	1.0	19
Union set goals (binary)	0.8	0.4	0.0	1.0	18
Union Tenure	29.1	23.7	4.0	87.0	19
Union Tenure President	17.6	15.2	6.0	72.0	17
Firm Tenure President	46.5	37.2	12.0	145.0	18
Firm Tenure LL	40.6	30.4	13.0	119.0	19
Firm Tenure Union W	31.4	22.4	9.1	78.2	17
Firm Tenure Non Union W	22.2	22.1	4.4	95.1	16
Sector Tenure President	76.4	64.0	20.0	246.0	18
Sector Tenure LL	72.8	44.9	25.8	167.6	19
Sector Tenure Union W	50.4	27.5	20.4	116.1	17
Sector Tenure Non Union W	46.3	29.9	16.6	142.8	16

Notes. Unit of observation is factory. The data in this table comes from the pre-sessions held by CTUM with the unions to explain about the intervention. The number of observations can be less than 19 factories as not all the factories had available the information requested. *Union set goals* is an indicator for whether the union has a stated goal. *Union Tenure* is number of months the union has been active at the factory. *Firm Tenure* indicates tenure at the factory (months) while *Sector Tenure* indicates tenure in the garment sector (months).

Table 2.7: Differences between Leaders and Workers, with controls

	Observations (1)	Worker Mean (2)	Coeff. on Line Leader (3)	Coeff. on President (4)	*p*-value of diff, cols (3)-(4) (5)
Panel A: Demographics & Ability					
Female	1103	.967	-0.137	-0.564	0.004
			[0.000]	[0.001]	
Age	1103	25.005	0.029	2.691	0.074
			[0.961]	[0.077]	
Migrant	1103	.52	-0.011	0.061	0.574
			[0.854]	[0.647]	
Education(Yrs)	1103	7.754	0.147	0.785	0.462
			[0.604]	[0.380]	
Raven Score	1103	4.524	0.126	1.637	0.022
			[0.680]	[0.021]	
Panel B: Employment & Minimum Wage Views					
Months in Factory	1103	29.888	2.646	5.320	0.645
			[0.390]	[0.406]	
Months in Sector	1103	50.621	13.540	8.939	0.615
			[0.000]	[0.353]	
Income	1103	243154.007	-1927.443	6825.805	0.660
			[0.619]	[0.734]	
Preferred Min Wage	1103	7504.258	73.261	199.725	0.628
			[0.647]	[0.517]	
Expected Min Wage	1103	6545.961	-209.817	-424.591	0.330
			[0.029]	[0.078]	
Panel C: Personality Traits					
Altruism	1103	1268.777	206.231	323.952	0.322
			[0.000]	[0.019]	
Extraversion	1103	3.392	0.230	0.481	0.244
			[0.009]	[0.038]	
Agreeableness	1103	3.862	0.319	0.102	0.442
			[0.000]	[0.701]	
Conscientiousness	1103	3.979	0.328	0.565	0.142
			[0.000]	[0.003]	
Neuroticism	1103	2.665	-0.277	-0.729	0.063
			[0.004]	[0.006]	
Openness	1103	3.001	-0.068	-0.426	0.064
			[0.436]	[0.033]	
Grit	1103	2.571	0.930	1.424	0.003
			[0.000]	[0.000]	
Locus of Control	1103	4.008	0.188	0.366	0.464
			[0.205]	[0.186]	
BFI Index	1103	2.314	0.341	0.468	0.728
			[0.637]	[0.902]	
Grit	1103	2.571	0.930	1.424	0.005
			[0.000]	[0.000]	
Locus of Control	1103	4.008	0.188	0.366	0.468
			[0.202]	[0.195]	

Notes: Unit of observation is worker. Probability weights used. All regressions include factory FE. With the exception of the BFI Index, each regression controls for all other variables included in the table. The BFI Index regression controls for all non-BFI variables in the table. *p*-values calculated using the wild cluster bootstrap-t method.

Table 2.8: Balance table: Consensus-building experiment

	(1)	(2)	(3)	(4)	(5)
		Mean / (SE)		**Difference in means / (p-value)**	
Variable	Control	Own LL	External LL	Diff Own-Control	Diff External-Control
Gender	1.022	1.033	1.061	0.005	0.025
	(0.148)	(0.178)	(0.239)	(0.659)	(0.160)
Age	25.737	23.929	24.552	-1.494***	-1.129**
	(6.440)	(5.556)	(5.792)	(0.000)	(0.037)
Education (Yrs)	7.627	7.969	7.675	0.327	-0.031
	(2.660)	(2.855)	(2.740)	(0.140)	(0.895)
Literacy	2.071	2.083	2.113	0.012	0.039
	(0.330)	(0.349)	(0.411)	(0.629)	(0.199)
Raven Score	4.376	4.895	4.654	0.457**	0.318
	(2.763)	(2.806)	(2.746)	(0.033)	(0.234)
Months in Factory	29.840	27.547	29.747	-0.521	0.150
	(33.458)	(30.497)	(36.326)	(0.801)	(0.943)
Months in Sector	52.257	42.634	50.913	-6.076**	2.010
	(50.759)	(43.124)	(53.266)	(0.038)	(0.626)
Min. Wage Belief	6,559.065	6,379.549	6,419.871	-114.294	-29.482
	(994.636)	(1,049.948)	(1,009.601)	(0.122)	(0.677)
Min. Wage Preference	7,523.598	7,248.997	7,295.476	-187.479	-116.892
	(1,557.759)	(1,514.251)	(1,540.256)	(0.108)	(0.350)
Absolute diff, worker and leader MW Preference	1,270.471	1,239.167	1,226.344	-66.739	-53.826
	(924.990)	(855.178)	(871.732)	(0.355)	(0.466)
Absolute diff, worker and leader MW Belief	776.069	799.399	924.622	-30.439	137.891**
	(639.849)	(634.670)	(707.401)	(0.577)	(0.038)
Grade	2.477	2.733	2.662	0.042	-0.110
	(1.403)	(1.416)	(1.479)	(0.563)	(0.235)
Last Month Income	242720.156	234366.094	234317.453	-3,114.145	-1,774.809
	(39,172.082)	(38,648.496)	(37,231.320)	(0.153)	(0.448)
Observations	425	284	206	709	631

Notes. Probability weights and standard errors clustered at the group level used. Controlling for factory FE x union status.

Table 2.9: Workers' awareness of a leader's participation in the group discussion

	Was there a LL in your discussion group?				
	(1)	(2)	(3)	(4)	(5)
Leader	0.409***				
	(0.0523)				
External Leader		0.222***			
		(0.0642)			
Own Leader		0.523***			
		(0.0574)			
External Leader, Union			0.188**		
			(0.0768)		
Own Leader, Union			0.549***		
			(0.0672)		
External Leader, Non-Union			0.280***		
			(0.0831)		
Own Leader, Non-Union			0.473***		
			(0.0721)		
Leader, High Similarity				0.323***	
				(0.0626)	
Leader, Low Similarity				0.487***	
				(0.0616)	
Own Leader, High Similarity					0.480***
					(0.0696)
External Leader, High Similarity					0.122*
					(0.0702)
Own Leader, Low Similarity					0.557***
					(0.0690)
External Leader, Low Similarity					0.335***
					(0.0868)
R-squared	0.2828	0.3288	0.3323	0.2968	0.3403
Control Mean	0.215	0.215	0.215	0.215	0.215
Number of obs.	746	746	746	746	746
p-values					
External = Own:		0.000			
External, Union = Own, Union:			0.000		
External, Non-Union = Own, Non-Union:			0.047		
External, Union =External, Non-Union:			0.344		
Own, Union = Own, Non-Union:			0.337		
High Similarity= Low Similarity:				0.013	
Own High Similarity= Ext High Similarity:					0.000
Own High Similarity= Own Low Similarity:					0.307
Ext High Similarity= Ext Low Similarity:					0.026
Own Low Similarity= Ext Low Similarity:					0.027

Notes. Unit of observation is worker. Probability weights used and standard errors clustered at the group level. Dependent variable is *LLInGroup*, the workers' belief about the presence of a union line leader or an EC member in their group. Stratification FEs are included: Factory FEs x Union FEs. Controlling for group size FEs. The sample size in this regression is smaller than the full worker sample (n=914) because 18% of workers incorrectly reported that they were line leaders in the follow-up survey and were not asked this question. In the Supplementary Materials, we report balance tests for the subset of workers with non-missing data for this question.

Table 2.10: Engagement in Group Discussions

	Enjoyment				Achievement of Consensus				Participation			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Leader	0.0942*				0.293***				0.0857			
	(0.0511)				(0.0782)				(0.0671)			
External Leader		0.0527				0.234**				0.140		
		(0.0657)				(0.114)				(0.0992)		
Own Leader		0.121**				0.331***				0.0503		
		(0.0558)				(0.0837)				(0.0701)		
Leader, High Similarity			0.100				0.287***				0.165**	
			(0.0610)				(0.0921)				(0.0836)	
Leader, Low Similarity			0.0884				0.299***				0.00757	
			(0.0601)				(0.0953)				(0.0787)	
Own Leader, High Similarity				0.137*				0.247**				0.0952
				(0.0699)				(0.0971)				(0.0839)
Own Leader, Low Similarity				0.107				0.401***				0.00890
				(0.0660)				(0.107)				(0.0890)
External Leader, High Similarity				0.0516				0.331**				0.255*
				(0.0857)				(0.138)				(0.134)
External Leader, Low Similarity				0.0529				0.134				0.0130
				(0.0842)				(0.154)				(0.110)
R-squared	0.062	0.064	0.062	0.064	0.099	0.100	0.099	0.105	0.070	0.072	0.075	0.078
Control Mean	0.007	0.007	0.007	0.007	0.000	0.000	0.000	0.000	-0.000	-0.000	-0.000	-0.000
Number of obs.	914	914	914	914	914	914	914	914	914	914	914	914
p-values												
External = Own:		0.290				0.403				0.365		
High Similarity= Low Similarity:			0.857				0.906				0.092	
Own High Similarity= Ext High Similarity:				0.375				0.565				0.262
Own High Similarity= Own Low Similarity:				0.703				0.214				0.405
Ext High Similarity= Ext Low Similarity:				0.990				0.291				0.117
Own Low Similarity= Ext Low Similarity:				0.530				0.127				0.972
Own High Similarity= Ext Low Similarity:				0.392				0.488				0.488

Notes. Unit of observation is worker in all columns. All three outcome variables are indexes of the following self-reported survey measures of participants' engagement. *Enjoyment* includes interest and enjoyment of the discussion as well whether the respondent perceived it to be worthwhile (*Group Interested*, *Group Enjoy*, *Group Unease*[reverse]), and *Group Waste*[reverse]. *Agreement* includes group consensus on minimum wage preferences and prediction (*Group Agree Ideal* and *Group Agree Prediction*). *Participation* includes freedom to express views(*Group Express Ideas*), and active participation by all members (*Group All Participate*). Probability weights and standard errors clustered at the group level. Controlling for group size FE and stratification FEs(Factory FEs x UnionFEs).

Table 2.11: Unannounced survey attendance results

	Attendance Survey						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Leader	0.114		0.114				
	(0.0981)		(0.0868)				
Own Leader		0.114		0.114			
		(0.109)		(0.0954)			
External Leader		0.113		0.113			
		(0.112)		(0.0979)			
Leader, High Similarity					0.135		
					(0.0941)		
Leader, Low Similarity					0.0788		
					(0.123)		
Leader, Union						0.00904	
						(0.130)	
Leader, Non-Union						0.319**	
						(0.129)	
Own Leader, High Similarity							0.153
							(0.115)
Own Leader, Low Similarity							0.0441
							(0.177)
External Leader, High Similarity							0.107
							(0.117)
External Leader, Low Similarity							0.127
							(0.146)
R-squared	0.415	0.415	0.391	0.391	0.412	0.431	0.418
Control Mean	0.341	0.341	0.341	0.341	0.341		0.341
Control Mean Union						0.320	
Control Mean Non-Union						0.385	
Number of obs.	117	117	117	117	117	117	117
p-values							
Own Leader = External Leader		0.995		0.994			
High Similarity= Low Similarity					0.650		
Leader Union = Leader Non-Union						0.089	
Own High Similarity= Ext High Similarity:							0.654
Own High Similarity= Own Low Similarity:							0.557
Ext High Similarity= Ext Low Similarity:							0.890
Own Low Similarity= Ext Low Similarity:							0.673
Own High Similarity= Ext Low Similarity:							0.862
PDS lasso selected controls	N	N	Y	Y	N	N	N

Notes. Unit of observation is worker in all columns. We only keep the workers that are in the control group in the mobilization experiment. No controls are selected for Col.3 and Col.4. R-squared for columns that applied PDS lasso selected controls are estimated by the correlation between the observed outcome and the predicted outcome. Probability weights and standard errors clustered at the group level. Controlling for group size FE and stratification FEs (Factory FEs x Union FEs).

Table 2.12: Group behavior, as assessed by research staff

	Share engaged (1)	Share distracted (2)	Active facilitation (3)	Asking opinions (4)	Summarizing opinions (5)	Taking notes (6)
Panel A: Leader						
Leader	0.0262	-0.0681**	-0.00180	-0.0135	0.173***	0.184***
	(0.0267)	(0.0277)	(0.0614)	(0.0616)	(0.0560)	(0.0550)
R-squared	0.176	0.202	0.211	0.304	0.331	0.299
Panel B: Own vs. External Leader						
Own Leader	0.0424	-0.110***	-0.0164	0.0327	0.186***	0.141**
	(0.0290)	(0.0281)	(0.0688)	(0.0681)	(0.0670)	(0.0594)
External Leader	0.00102	-0.00365	0.0210	-0.0852	0.153**	0.249***
	(0.0351)	(0.0362)	(0.0819)	(0.0814)	(0.0702)	(0.0786)
R-squared	0.184	0.247	0.212	0.314	0.331	0.309
p-values						
Own Leader = External Leader	0.233	0.001	0.660	0.150	0.683	0.181
Panel C: High vs. Low Similarity Leader						
Leader Group, High Similarity (50th)	-0.00196	-0.0572*	0.0313	0.0752	0.186***	0.245***
	(0.0314)	(0.0312)	(0.0692)	(0.0724)	(0.0647)	(0.0634)
Leader Group, Low Similarity	0.0537*	-0.0787**	-0.0341	-0.100	0.161**	0.124*
	(0.0304)	(0.0325)	(0.0789)	(0.0722)	(0.0702)	(0.0684)
Stratification FEs	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.190	0.204	0.214	0.326	0.331	0.311
Control Group Mean	0.819	0.203	0.721	0.464	0.264	0.654
Number of obs.	201	201	201	201	201	201
p-values						
High Similarity = Low Similarity	0.084	0.499	0.430	0.025	0.743	0.112

Notes. Unit of observation is discussion group. Probability weights and robust standard errors used. Dependent variables are: *ShareEngaged*, the share of workers within a group that are engaged in the discussion; *ShareDistracted*, the share of workers within a group that are distracted during the discussion; *ActiveFacilitation*, an indicator for whether someone is actively facilitating the group; *AskingOpinions*, an indicator for whether someone is active others' opinions; *SummarizingOpinions*, an indicator for whether someone is summarizing opinions in the group; *TakingNotes*, an indicator for whether someone is taking notes in the group. Two members of the field team rated each group, and we average their observations in the analysis. Stratification FEs are Factory FEs. Controlling for group size FEs.

Table 2.13: Number of words in group discussion transcript

	Log(Total Word Count)				Log(Likely Worker Word Count)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Leader	-0.139				-0.401***			
	(0.0940)				(0.124)			
Own Leader		-0.263**				-0.531***		
		(0.111)				(0.137)		
External Leader		0.0325				-0.174		
		(0.111)				(0.140)		
Leader, High Similarity			-0.0623				-0.291**	
			(0.103)				(0.131)	
Leader, Low Similarity			-0.235*				-0.539***	
			(0.120)				(0.152)	
Own Leader, High Similarity				-0.141				-0.309**
				(0.123)				(0.147)
Own Leader, Low Similarity				-0.424***				-0.807***
				(0.156)				(0.181)
External Leader, High Similarity				0.0751				-0.171
				(0.126)				(0.147)
External Leader, Low Similarity				-0.00932				-0.173
				(0.145)				(0.184)
R-squared	0.236	0.271	0.248	0.290	0.372	0.406	0.388	0.442
Control Mean	1005	1005	1005	1005	876	876	876	876
Number of obs.	167	167	167	167	167	167	167	167
p-values								
Own Leader= External Leader:		0.015				0.011		
High Similarity= Low Similarity:			0.159				0.069	
Own High Similarity= Own Low Similarity:				0.098				0.006
Ext High Similarity= Ext Low Similarity:				0.591				0.990
Own High Similarity= Ext High Similarity:				0.137				0.314
Own High Similarity= Ext Low Similarity:				0.394				0.424
Own Low Similarity= Ext High Similarity:				0.005				0.002
Own Low Similarity= Ext Low Similarity:				0.034				0.006

Notes. Unit of observation is discussion group. Probability weights and robust standard errors used. Dependent variables are: *Log(Total Word Count)*, the number of words spoken by the group members; *Log(Likely Worker Word Count)*, the number of words spoken by possible workers (= group members who are not identified as a confirmed/possible leader). Stratification FEs are Factory FEs. Columns (1)–(4) additionally control for the fixed effect of the number of group members. The number of possible workers is systematically lower in Leader groups because leaders are more likely to be identified in these groups. This mechanically reduces the number of words by possible workers in control groups. Therefore, in columns (5)–(8), we control for the fixed effects of the number of possible workers. Control mean shows the average number of words in control group before taking a logarithm.

Table 2.14: Word counts of responses to question prompts in group discussion experiment

	Log(Total Words)			Log(Topic 1: Benefit)			Log(Topic 2: Harm)			Log(Topic 3: Heterogeneity)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Leader Group	0.211**			0.203			0.183			0.231		
	(0.106)			(0.128)			(0.139)			(0.161)		
Own Leader Group		0.272**			0.241*			0.180			0.489***	
		(0.110)			(0.133)			(0.145)			(0.167)	
External Leader Group		0.116			0.144			0.188			-0.170	
		(0.124)			(0.154)			(0.168)			(0.216)	
Leader Group, High Similarity			0.219*			0.220			0.195			0.211
			(0.114)			(0.140)			(0.148)			(0.182)
Leader Group, Low Similarity			0.202*			0.184			0.170			0.254
			(0.120)			(0.144)			(0.162)			(0.198)
R-squared	0.178	0.191	0.178	0.092	0.095	0.092	0.127	0.127	0.127	0.187	0.244	0.187
Control Mean	36.949	36.949	36.949	11.951	11.951	11.951	13.580	13.580	13.580	11.419	11.419	11.419
Number of obs.	202	202	202	202	202	202	202	202	202	202	202	202
p-values												
P-val: Own LL= External LL		0.110			0.433			0.958			0.002	
P-val: High Quality=Low Quality			0.861			0.767			0.855			0.828

Notes. Unit of observation is discussion group. Probability weights and robust standard errors used. Dependent variables are the number of words written in group discussion form as the group's answers to each of the following questions: (1) How do you think that a minimum wage increase may benefit workers? (*Topic 1: Benefit*); (2) How do you think that a minimum wage increase may harm workers? (*Topic 2: Harm*); (3) Do you think it will affect different groups of workers, for example, skilled versus unskilled, union members versus non-members, differently? (*Topic 3: Heterogeneity*). Stratification FEs are Factory FEs. Fixed effect of the number of group members is also controlled. Control mean shows the average number of words in control group before taking a logarithm.

Table 2.15: Leader Behavior without control group, as assessed by research staff

	Speaking	Listening	Consensus building	Conflict resolution	Leadership
	(1)	(2)	(3)	(4)	(5)
Panel A: Summary statistics					
Mean Dependent Var.	4.6	4.7	4.1	3.3	4.5
S.D. Dependent Var.	1.6	1.4	1.7	1.8	1.6
Panel B: Own vs. External Leader					
Own Leader	0.225	-0.529	-0.260	-0.560	0.225
	(0.371)	(0.326)	(0.362)	(0.365)	(0.414)
Stratification FEs	Yes	Yes	Yes	Yes	Yes
R-squared	0.185	0.244	0.364	0.368	0.196
External Leader Group Mean	4.437	4.825	3.941	3.179	4.272
Number of obs.	119	119	119	119	119
Panel C: High vs. Low Similarity Leader, with factory (stratification) FE					
High Similarity Leader	0.798**	0.425	0.576*	0.329	0.525
	(0.332)	(0.298)	(0.301)	(0.379)	(0.351)
Stratification FEs	Yes	Yes	Yes	Yes	Yes
R-squared	0.231	0.235	0.382	0.357	0.212
Low Similarity Leader Group Mean	4.175	4.400	3.632	2.897	4.219
Number of obs.	119	119	119	119	119
Panel D: High vs. Low Similarity Leader, without factory (stratification) FE					
High Similarity Leader	0.927***	0.627**	0.809**	0.594	0.708**
	(0.304)	(0.279)	(0.339)	(0.402)	(0.339)
Stratification FEs	No	No	No	No	No
R-squared	0.083	0.082	0.066	0.039	0.052
Low Similarity Leader Group Mean	4.175	4.400	3.632	2.897	4.219
Number of obs.	119	119	119	119	119

Notes. Unit of observation is discussion group. Probability weights and robust standard errors used. The dependent variables are: *Speaking*, assessing the extent of LL speaking; *Listening*, assessing the extent of LL listening; *ConsensusBuilding*, assessing the extent of LL engaged in consesus building; *ConflictResolution*, assessing the extent of LL engaged in conflict resolution; and *Leadership*, assessing the extent of LL showing leadership. All dependent variables are measured on a Likert scale 1-7 separately by two members of the research staff and the average is taken. In Panels B and C, Stratification FEs are Factory FEs. Controlling for group size FEs.

Table 2.16: Leader similarity and group discussion results

	Deviation from Union Preference	Deviation from Union Belief	Engagement	Active Group
	(1)	(2)	(3)	(4)
Panel A: High vs. Low Similarity Leaders				
Leader, High Similarity	-271.3**	-31.26	0.161***	0.191***
	(107.6)	(69.11)	(0.0565)	(0.0617)
Leader, Low Similarity	-140.6	-15.22	0.113**	0.0497
	(106.9)	(67.81)	(0.0525)	(0.0687)
R-squared	0.232	0.335	0.084	0.400
Control Mean	1130.078	712.308	-0.039	0.127
Number of obs.	914	914	914	202
p-values				
High Similarity= Low Similarity:	0.213	0.849	0.444	0.074
Panel B: High vs. Low Similarity Leaders, Own vs. External				
Own Leader, High Similarity	-262.6**	-25.74	0.148**	0.155**
	(117.5)	(80.47)	(0.0578)	(0.0730)
Own Leader, Low Similarity	-244.4*	-38.11	0.142**	0.0856
	(130.8)	(83.68)	(0.0596)	(0.0858)
External Leader, High Similarity	-278.1*	-36.94	0.175*	0.233***
	(147.1)	(106.0)	(0.0897)	(0.0852)
External Leader, Low Similarity	16.55	19.73	0.0692	-0.00355
	(132.8)	(85.83)	(0.0763)	(0.0999)
R-squared	0.2386	0.3362	0.0859	0.4055
Control Mean	1130.078	712.308	-0.039	0.127
Number of obs.	914	914	914	202
p-values				
Own High Similarity= Ext High Similarity:	0.919	0.928	0.772	0.427
Own High Similarity= Own Low Similarity:	0.896	0.906	0.930	0.511
Ext High Similarity= Ext Low Similarity:	0.077	0.658	0.336	0.044
Own Low Similarity= Ext Low Similarity:	0.098	0.578	0.390	0.471

Notes. Unit of observation is worker in all columns but in Col. 4, where it is discussion group. The variable *Leader, High Similarity* is a binary variable equal to 1 if the estimated probability of a line leader having similar attributes to president is above the median. The probabilities are estimated for each worker based on a probit model, which includes demographics (gender, age, education, migrant(0/1), months in factory/sector), personality metrics (extraversion, agreeableness, conscientionsness, neuroticism, openness) and psychological metrics (raven, score, grit, altruism, choice in life).*Engagement* is an index of the following self-reported survey measures of participants' engagement: group consensus on minimum wage prediction/preferences (*GroupAgreePrediction*, *GroupAgreeIdeal*), freedom to express views (*GroupExpressIdeas*, *GroupUnease*[reverse]), interest and enjoyment of the discussion (*GroupInterested*, *GroupEnjoy*, *GroupWaste*[reverse]) and active participation by all members (*GroupAllParticipate*). *Active Group* is an index created from research staff observations which assess group behavior (*ShareEngaged, ShareDistracted, ActiveFacilitation AskingOpinions SummerizingOpinions TakingNotes*). The dependent variables in col. 1 and 2 represent the deviation from the factory average of baseline leaders' preferences and views respectively. Probability weights and standard errors clustered at the group level. Controlling for group size FE and stratification FEs (Factory FEs x Union FEs).

Table 2.17: Union affiliation and group discussion results

	Deviation from Union Preference	Deviation from Union Belief	Engagement
	(1)	(2)	(3)
Leader, Union	-277.5**	-9.027	0.110**
	(112.3)	(61.88)	(0.0540)
Leader, Non-Union	-109.7	-42.05	0.226***
	(120.2)	(68.66)	(0.0739)
R-squared	0.248	0.340	0.092
Control Mean Union	1205.288	712.767	-0.004
Control Mean Non-Union	995.156	711.485	-0.102
Number of obs.	914	914	914
p-values			
Leader Union = Leader Non-Union	0.238	0.638	0.174

Notes. Unit of observation is worker in all columns. *Engagement* is an index of the following self-reported survey measures of participants' engagement: group consensus on minimum wage prediction/preferences (*GroupAgreePrediction*, *GroupAgreeIdeal*), freedom to express views (*GroupExpressIdeas*, *GroupUnease*[reverse]), interest and enjoyment of the discussion (*GroupInterested*, *GroupEnjoy*, *GroupWaste*[reverse]) and active participation by all members (*GroupAllParticipate*). The dependent variables in col. 1 and 2 represent the deviation from the factory average of baseline leaders' views and preferences respectively. Probability weights and standard errors clustered at the group level. Controlling for group size FE and stratification FEs (Factory FEs x Union FEs).

Table 2.18: Placebo control group leaders, leader similarity, and main results, control group leader is the member with highest similarity

	Deviation from Union Pref.	Deviation Exc. Leader	Deviation from Union or Placebo leader	Deviation from Union Belief	Deviation Exc. Leader	Deviation from Union or Placebo leader	Engagement	Active Group
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Leader, High Similarity	-273.2**	-224.4*	-511.1***	-39.32	-40.76	98.92	0.161**	0.152**
	(128.7)	(132.3)	(181.7)	(91.83)	(93.44)	(103.0)	(0.0705)	(0.0689)
Leader, Low Similarity	-182.8	-140.1	-463.9**	-19.36	-15.05	83.86	0.0958	0.00902
	(129.4)	(132.7)	(189.3)	(92.52)	(93.82)	(106.8)	(0.0682)	(0.0770)
Control, High Similarity	23.95	5.078	-230.5	4.793	2.508	-25.69	0.0191	-0.0903
	(153.1)	(155.7)	(210.2)	(93.73)	(93.80)	(135.8)	(0.0807)	(0.0678)
R-squared	0.248	0.251	0.275	0.336	0.330	0.379	0.086	0.411
Control Mean	1135.491	1135.491	1395.362	713.572	713.572	684.124	-0.039	0.127
Number of obs.	833	833	833	833	833	833	833	202
p-values								
High Similarity= Low:	0.379	0.433	0.676	0.813	0.765	0.870	0.309	0.076
Leader High= Control High:	0.027	0.093	0.063	0.579	0.588	0.308	0.047	0.001
Leader Low= Control High:	0.147	0.317	0.214	0.752	0.818	0.423	0.291	0.232

Notes. Unit of observation is worker in all columns but in Col. 8, where it is discussion group. The variable *Leader Group, High Similarity* is a binary variable equal to 1 if the estimated probability of a line leader having similar attributes to president is above the median. In the control group, the worker with the highest probit scores is considered as the leader (placebo leader). Sample restricted to workers who are not placebo leaders. The probabilities are estimated for each worker based on a probit model which includes demographics (gender, age, education, migrant(0/1), months in factory/sector), personality metrics (extraversion, agreeableness, conscientionsness, neuroticism, openness) and psychological metrics (raven, score, grit, altruism, choice in life). The dependent variables in columns 1-6 represent the deviation from the factory average of baseline leaders' preferences and views (cols. 1 and 4, respectively); cols. 2 and 5 exclude the individual leader view from the factory average for the leader groups and cols. 3 and 6 use deviation from the placebo leader view for the control groups. *Engagement* is an index of the following self-reported survey measures of participants' engagement: group consensus on minimum wage prediction/preferences (*GroupAgreePrediction*, *GroupAgreeIdeal*), freedom to express views (*GroupExpressIdeas*, *GroupUnease*[reverse]), interest and enjoyment of the discussion (*GroupInterested*, *GroupEnjoy*, *GroupWaste*[reverse]) and active participation by all members (*GroupAllParticipate*). *Active Group* is an index created from research staff observations which assess group behavior (*ShareEngaged, ShareDistracted, ActiveFacilitation AskingOpinions SummerizingOpinions TakingNotes*). Probability weights and standard errors clustered at the group level. Controlling for group size FE and stratification FEs (Factory FEs x Union FEs).

Table 2.19: Balance table: Mobilization, Coordination, and Social Pressure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Difference in means / (p-value)						
Variable	LL	LL & Info Least	LL & Info Most	LL & Sanctioning	Sanctioning	Info Least	Info Most
Gender	-0.047	0.136	-0.138	-0.033	-0.012	-0.010	-0.000
	(0.567)	(0.273)	(0.242)	(0.697)	(0.756)	(0.759)	()
Age	-2.938**	3.277*	-0.001	-1.206	-0.050	0.488	10.000
	(0.039)	(0.085)	(1.000)	(0.329)	(0.967)	(0.696)	(0.226)
Education (Yrs)	-0.333	-0.143	0.398	-0.140	-0.065	-0.566	-2.000*
	(0.636)	(0.888)	(0.783)	(0.851)	(0.917)	(0.430)	(0.056)
Literacy	-0.005	0.211	-0.073	-0.043	-0.103	-0.075	-0.000
	(0.945)	(0.197)	(0.640)	(0.478)	(0.219)	(0.351)	()
Raven Score	-0.472	-0.798	0.331	0.690	-0.590	0.005	-3.000***
	(0.555)	(0.413)	(0.767)	(0.365)	(0.334)	(0.995)	(0.005)
Months in Factory	-5.990	8.601	16.928	2.884	-7.121	-4.760	1.500
	(0.292)	(0.528)	(0.170)	(0.492)	(0.111)	(0.400)	(0.889)
Months in Sector	-13.323	19.169	8.184	1.083	4.158	-1.715	13.500
	(0.160)	(0.221)	(0.654)	(0.888)	(0.595)	(0.860)	(0.558)
Min. Wage Belief	-326.645	-184.356	-178.740	-18.578	-105.972	106.890	-100.000
	(0.170)	(0.558)	(0.701)	(0.938)	(0.613)	(0.664)	(0.331)
Min. Wage Preference	138.246	-3.007	999.967	231.909	238.446	256.884	600.000
	(0.643)	(0.995)	(0.155)	(0.467)	(0.484)	(0.437)	(0.331)
Grade	0.129	-0.472	-0.151	0.115	-0.175	0.014	0.000
	(0.645)	(0.295)	(0.779)	(0.627)	(0.441)	(0.950)	()
Last Month Income	-12242.940	6,238.098	-1,156.007	-12952.256*	-6,105.475	-5,149.977	-9,000.000
	(0.222)	(0.518)	(0.914)	(0.082)	(0.423)	(0.215)	(0.381)
Observations	257	145	214	251	254	228	161

Notes. Probability weights and robust standard errors used. Controlling for factory FE x discussion group FE. Showing the difference in means and p-values in parenthesis.

2.13 Appendix: Placebo and robustness tests for consensus-building experiment

We conduct two placebo tests for the main results. For the first test, for each control discussion group, we identify the worker with the highest predicted leader similarity score, and we assign this worker as the placebo leader for the group. For leader groups, we use the assigned leader's baseline view. We test whether we identify greater convergence in treatment groups to the real leader's view compared to the placebo leader's view. Table 2.20 presents the results. Column (1) shows that we find much stronger convergence to the real leaders' minimum wage preferences relative to the placebo leaders' preferences. Column (2) shows that the evidence of convergence is especially strong for the own leader treatment arm, although there is suggestive evidence of greater convergence to the external leaders' preferences compared to the placebo control leaders'. Consistent with our main results, we find no evidence of effects on expectations about the likely minimum wage level (columns (4) and (5)).

Next, we return to the baseline construction of our outcome variable, but for groups assigned to the external leader arm, we test for convergence to the *internal* union leaders' average preferences (beliefs). If our main specification is simply picking up the fact that leaders and workers have different preferences (beliefs), so that we are capturing the effect of having *any* union leader participating, then we would expect to find a similar amount of convergence in the external arm using the internal leaders' preferences (beliefs). Table 2.20 presents the results. Column (3) shows that the estimated coefficient for the external leader arm is about 50% smaller compared to the estimate for this group in column (4) of Table 2.3; these coefficients are statistically different at the 5% level. Evidently, our main specification is not simply capturing convergence due to *any* leaders' participation; instead, it is capturing convergence to the position of the participating leader's union. Consistent with our main results, column (6) shows that there is no evidence of convergence in beliefs for the external arm we use own and when we use internal union leaders' beliefs.

Table 2.20: Placebo control group leaders, control group leader is the member with highest similarity, and replacing deviation from external leaders with deviation from own leaders' view

	Ideal: Predicted Leader Control		Ideal: Own views for External	Guess: Predicted Leader Control		Guess: Own views for External
	(1)	(2)	(3)	(4)	(5)	(6)
Leader	-256.7**			136.1		
	(122.0)			(83.68)		
Own Leader		-288.3**	-267.4**		127.6	-19.02
		(143.0)	(104.4)		(95.73)	(64.96)
External Leader		-206.8	-76.65		149.6	23.57
		(153.7)	(109.9)		(109.3)	(78.07)
R-squared	0.286	0.286	0.220	0.485	0.485	0.304
Control Mean	1395.362	1395.362	1130.078	684.124	684.124	712.308
Number of obs.	833	833	914	833	833	914
p-values						
External=Own:		0.628	0.065		0.850	0.612

Notes. Unit of observation is worker. Probability weights and standard errors clustered at the group level. Columns 1,2,4,5: For groups with leaders, the dependent variable is the absolute value of the endline minimum wage guess/ideal minus the leader baseline view; for control groups, the dependent variable is the absolute value of the endline minimum wage guess/ideal minus the worker of highest similarity baseline view (placebo leader); sample restricted to workers who are not placebo leaders. Columns 3 and 6: using main specification as in Table 3 but, for external LL groups, replacing the deviation from the external leaders view with deviation from own factory leaders view. The p-values when testing col. 3 coefficients with those in Table 3 col. 1 are: 0.32 for *Own Leader* and 0.05 for *External Leader*. The p-values when testing col. 6 coefficients with those in Table 3 col. 3 are: 0.79 for *Own Leader* and 0.67 for *External Leader*. Stratification FEs are included: Factory FEs x Union FEs. Controlling for group size FE.

We also conduct multiple robustness checks. First, we check whether union leaders have effects on group discussion outcomes even conditional on the predicted leader similarity of the workers in their discussion group. In Appendix Tables 2.21, 2.22, and 2.23 we show that our results hold controlling for the average (2.21) or the maximum (2.22) of the similarity score among workers in the discussion group; we also run a flexible specification where we rank group participants by their similarity score and control for the similarity of each rank (2.23). It is clear that leaders influence groups' outcomes above and beyond even other potentially prominent individuals in the group. Second, the results hold when controlling for the leader or placebo leader similarity (Appendix Table 2.24). Third, we conduct a robustness test for our leader similarity measure, which is that we drop one family of variables in the prediction model at a time (i.e., demographics, personality traits, psychological traits, and education/tenure) and re-estimate the results. Our results are robust to dropping each family of variables (results not reported). Fourth, we show in Table 2.25 that results hold if we do not use probability weights in the regressions.

Finally, leaders are somewhat more likely than workers to be men (12.9% compared to 3.3%). Gender is an observable characteristic that may be an alternative channel through which leaders affect workers or may complement or substitute for leadership. Consequently, we separately test for the effects of female and male leaders in Table 2.26. While the smaller sample of male leaders limits our precision for this group, the effects do not provide evidence of heterogeneity except for the deviation from the union's beliefs about the minimum wage; in addition to inducing convergence in preferences to the union's ideal (of a similar magnitude as female leaders), male leaders also induce convergence in beliefs to that of the union's leaders. As the leader's gender also affects the group's gender composition, which may affect consensus building through other channels than leadership, Table 2.27 shows that our main results are robust to controlling for groups' gender composition.

Table 2.21: Average discussion group leader similarity and union leader

	Deviation from Union Preference	Deviation from Union Belief	Engagement	Active Group
	(1)	(2)	(3)	(4)
Leader	-192.7***	-25.85	0.138***	0.130**
	(60.73)	(33.66)	(0.0411)	(0.0604)
Average Group Similarity	-3496.1**	702.4	1.192	0.0376
	(1433.9)	(1043.9)	(0.948)	(1.854)
R-squared	0.250	0.340	0.091	0.393
Control Mean	1130.078	712.308	-0.039	0.127
Number of obs.	914	914	914	202

Notes. Unit of observation is worker in all columns but in Col. 4, where it is discussion group. *Engagement* is an index of the following self-reported survey measures of participants' engagement: group consensus on minimum wage prediction/preferences (*GroupAgreePrediction*, *GroupAgreeIdeal*), freedom to express views (*GroupExpressIdeas*, *GroupUnease*[reverse]), interest and enjoyment of the discussion (*GroupInterested*, *GroupEnjoy*, *GroupWaste*[reverse]) and active participation by all members (*GroupAllParticipate*). *Active Group* is an index created from research staff observations which assess group behavior (*ShareEngaged, ShareDistracted, ActiveFacilitation AskingOpinions SummerizingOpinions TakingNotes*). The dependent variables in col. 1 and 2 represent the deviation from the factory average of baseline leaders' preferences and views respectively. Probability weights and bootstrap standard errors clustered at the group level. Controlling for group size FE and stratification FEs (Factory FEs x Union FEs).

Table 2.22: Maximum discussion group leader similarity and union leader

	Deviation from Union Preference	Deviation from Union Belief	Engagement	Active Group
	(1)	(2)	(3)	(4)
Leader	-196.4***	-26.09	0.137***	0.130**
	(60.58)	(33.54)	(0.0411)	(0.0605)
Max Similarity in Group	-638.6**	152.6	0.260	0.0115
	(286.2)	(216.9)	(0.189)	(0.381)
R-squared	0.249	0.340	0.091	0.393
Control Mean	1130.078	712.308	-0.039	0.127
Number of obs.	914	914	914	202

Notes. Unit of observation is worker in all columns but in Col. 4, where it is discussion group. *Engagement* is an index of the following self-reported survey measures of participants' engagement: group consensus on minimum wage prediction/preferences (*GroupAgreePrediction*, *GroupAgreeIdeal*), freedom to express views (*GroupExpressIdeas*, *GroupUnease*[reverse]), interest and enjoyment of the discussion (*GroupInterested*, *GroupEnjoy*, *GroupWaste*[reverse]) and active participation by all members (*GroupAllParticipate*). *Active Group* is an index created from research staff observations which assess group behavior (*ShareEngaged, ShareDistracted, ActiveFacilitation AskingOpinions SummerizingOpinions TakingNotes*). The dependent variables in col. 1 and 2 represent the deviation from the factory average of baseline leaders' preferences and views respectively. Probability weights and bootstrap standard errors clustered at the group level. Controlling for group size FE and stratification FEs (Factory FEs x Union FEs).

Table 2.23: Robustness to flexibly controlling for group similarity

	Deviation from Union Preference	Deviation from Union Belief	Engagement	Active Group
	(1)	(2)	(3)	(4)
Leader	-170.1***	-24.16	0.130***	0.129**
	(59.57)	(34.37)	(0.0414)	(0.0598)
Similarity of Member w/ Rank=1	-479.0*	170.0	0.281	0.0134
	(287.9)	(237.5)	(0.193)	(0.425)
Similarity of Member w/ Rank=2	7824.0**	45.08	-2.712	4.292
	(3509.8)	(2242.0)	(2.263)	(3.548)
Similarity of Member w/ Rank=3	-45162.9**	-793.5	-2.733	-39.79
	(18838.8)	(14404.0)	(16.32)	(29.85)
Similarity of Member w/ Rank=4	-143344.6**	-15337.1	77.27	100.5
	(62625.8)	(51537.7)	(61.55)	(122.1)
R-squared	0.269	0.340	0.096	0.404
Control Mean	1130.078	712.308	-0.039	0.127
Number of obs.	914	914	914	202

Notes. Unit of observation is worker in all columns but in Col. 4, where it is discussion group. *Engagement* is an index of the following self-reported survey measures of participants' engagement: group consensus on minimum wage prediction/preferences (*GroupAgreePrediction*, *GroupAgreeIdeal*), freedom to express views (*GroupExpressIdeas*, *GroupUnease*[reverse]), interest and enjoyment of the discussion (*GroupInterested*, *GroupEnjoy*, *GroupWaste*[reverse]) and active participation by all members (*GroupAllParticipate*). *Active Group* is an index created from research staff observations which assess group behavior (*ShareEngaged, ShareDistracted, ActiveFacilitation AskingOpinions SummerizingOpinions TakingNotes*). The dependent variables in col. 1 and 2 represent the deviation from the factory average of baseline leaders' preferences and views respectively. Probability weights and bootstrap standard errors clustered at the group level. Controlling for group size FE and stratification FEs (Factory FEs x Union FEs).

Table 2.24: Robustness to controlling for leader & placebo leader similarity

	Deviation from Union Preference	Deviation from Union Belief	Engagement	Active Group
	(1)	(2)	(3)	(4)
Leader	-218.0***	-27.35	0.149***	0.126**
	(59.05)	(32.84)	(0.0404)	(0.0596)
Leader or placebo leader similarity	-174.2	267.2	-0.0267	0.180
	(292.1)	(224.1)	(0.174)	(0.422)
R-squared	0.246	0.341	0.090	0.394
Control Mean	1130.078	712.308	-0.039	0.127
Number of obs.	914	914	914	202

Notes. Unit of observation is worker in all columns but in Col. 4, where it is discussion group. *Engagement* is an index of the following self-reported survey measures of participants' engagement: group consensus on minimum wage prediction/preferences (*GroupAgreePrediction*, *GroupAgreeIdeal*), freedom to express views (*GroupExpressIdeas*, *GroupUnease*[reverse]), interest and enjoyment of the discussion (*GroupInterested*, *GroupEnjoy*, *GroupWaste*[reverse]) and active participation by all members (*GroupAllParticipate*). *Active Group* is an index created from research staff observations which assess group behavior (*ShareEngaged, ShareDistracted, ActiveFacilitation AskingOpinions SummerizingOpinions TakingNotes*). The dependent variables in col. 1 and 2 represent the deviation from the factory average of baseline leaders' preferences and views respectively. Probability weights and bootstrap standard errors clustered at the group level. Controlling for group size FE and stratification FEs (Factory FEs x Union FEs).

Table 2.25: Robustness to estimating consensus-building results without probability weights

	Deviation from Union Preference		Deviation from Union Belief		Engagement		Active Group	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Leader								
Leader	-196.6**	-195.3**	-2.145	-9.246	0.115**	0.113**	0.147***	0.145***
	(85.67)	(84.54)	(57.61)	(55.74)	(0.0487)	(0.0449)	(0.0531)	(0.0507)
R-squared	0.223	0.229	0.296	0.307	0.069	0.140	0.395	0.403
Panel B: Own versus External LL								
External Leader	-153.5	-155.3	29.36	13.78	0.0788	0.0855	0.128	0.119
	(99.77)	(98.27)	(76.89)	(75.06)	(0.0676)	(0.0614)	(0.0777)	(0.0745)
Own Leader	-224.1**	-220.7**	-21.84	-23.42	0.139***	0.131***	0.160***	0.161***
	(98.08)	(96.70)	(67.84)	(65.33)	(0.0509)	(0.0469)	(0.0572)	(0.0541)
R-squared	0.224	0.230	0.297	0.307	0.070	0.141	0.396	0.404
Control Mean	1132.922	1132.922	719.659	719.659	0.002	0.002	0.098	0.098
Number of obs.	914	914	914	914	914	914	202	202
p-values								
Own Leader = External Leader	0.486	0.509	0.557	0.660	0.358	0.434	0.691	0.576
PDS lasso selected controls	N	Y	N	Y	N	Y	N	Y

Notes. Unit of observation is worker in all columns but in Col. 7 and Col.8, where it is discussion group.*Engagement* is an index of the following self-reported survey measures of participants' engagement: group consensus on minimum wage prediction/preferences (*GroupAgreePrediction*, *GroupAgreeIdeal*), freedom to express views (*GroupExpressIdeas*, *GroupUnease*[reverse]), interest and enjoyment of the discussion (*GroupInterested*, *GroupEnjoy*, *GroupWaste*[reverse]) and active participation by all members (*GroupAllParticipate*). *Active Group* is an index created from research staff observations which assess group behavior (*ShareEngaged, ShareDistracted, ActiveFacilitation AskingOpinions SummerizingOpinions TakingNotes*). The dependent variables in col. 1, col.2 and col.3, col.4 represent the deviation from the factory average of baseline leaders' preferences and views respectively. Standard errors clustered at the group level. Controlling for group size FE and stratification FEs (Factory FEs x Union FEs). Selected controls for Col.6 are *Grit* and *Agreeableness squared (BFI)*. *Deviation from Union Preference* is selected for Col.2. No controls are selected for Col.4 and Col.8. R-squared for columns that applied PDS lasso selected controls are estimated by the correlation between the observed outcome and the predicted outcome.

Table 2.26: Testing for heterogeneous treatment effects by gender of leader

	Deviation from Union Preference		Deviation from Union Belief		Engagement		Group Activity	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Leader								
FemaleLeader	-158.2	-159.2	22.73	8.887	0.134***	0.121**	0.121**	0.123**
	(105.9)	(104.0)	(61.98)	(59.76)	(0.0514)	(0.0474)	(0.0600)	(0.0571)
MaleLeader	-381.6***	-375.5***	-124.8	-117.2*	0.182***	0.180***	0.156*	0.142
	(131.5)	(123.5)	(77.15)	(69.98)	(0.0643)	(0.0583)	(0.0924)	(0.0885)
R-squared	0.253	0.240	0.347	0.314	0.091	0.126	0.394	0.401
Panel B: Own versus External LL								
ExternalFemaleLeader	-71.77	-74.94	49.54	30.07	0.112	0.109	0.133*	0.136*
	(123.6)	(121.2)	(96.96)	(94.23)	(0.0748)	(0.0681)	(0.0777)	(0.0737)
ExternalMaleLeader	-444.8**	-439.1**	-150.1*	-145.6*	0.164	0.156	0.151	0.111
	(198.4)	(185.2)	(76.51)	(76.16)	(0.116)	(0.100)	(0.150)	(0.149)
OwnFemaleLeader	-216.7*	-215.5*	4.852	-5.197	0.148***	0.128**	0.114	0.114
	(118.6)	(116.9)	(69.98)	(66.93)	(0.0560)	(0.0522)	(0.0756)	(0.0714)
OwnMaleLeader	-329.7**	-326.7**	-104.3	-95.79	0.195***	0.197***	0.161	0.165*
	(152.1)	(147.5)	(113.5)	(102.1)	(0.0659)	(0.0639)	(0.103)	(0.0958)
R-squared	0.256	0.242	0.347	0.315	0.091	0.127	0.394	0.402
Control Mean	1130.078	1130.078	712.308	712.308	-0.039	-0.039	0.127	0.127
Number of obs.	914	914	914	914	914	914	202	202
PDS lasso selected controls	N	Y	N	Y	N	Y	N	Y

Notes. Unit of observation is worker in all columns but in Col. 7 and Col.8, where it is discussion group.

Table 2.27: Robustness to controlling for the share of men in the group discussion

	Deviation from Union Preference		Deviation from Union Belief		Self-reported Engagement		Observed Group Activity	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Leader								
Leader	-180.0*	-178.4*	15.20	5.314	0.110**	0.0996**	0.237***	0.239***
	(96.86)	(95.29)	(55.31)	(53.32)	(0.0450)	(0.0419)	(0.0829)	(0.0791)
R-squared	0.257	0.242	0.356	0.324	0.081	0.131	0.352	0.352
Panel B: Own versus External LL								
External Leader	-135.3	-136.1	28.45	12.87	0.0883	0.0844	0.190*	0.184*
	(110.9)	(108.8)	(79.90)	(77.75)	(0.0610)	(0.0556)	(0.111)	(0.106)
Own Leader	-208.4*	-205.5*	6.926	0.630	0.123**	0.109**	0.267***	0.274***
	(109.0)	(107.2)	(63.38)	(60.71)	(0.0491)	(0.0454)	(0.0912)	(0.0867)
R-squared	0.258	0.243	0.356	0.325	0.081	0.132	0.355	0.361
Control Mean	1130.078	1130.078	712.308	712.308	-0.090	-0.090	-0.090	-0.090
Number of obs.	914	914	914	914	914	914	202	202
p-values								
Own Leader = External Leader	0.492	0.504	0.809	0.887	0.571	0.647	0.485	0.396
PDS lasso selected controls	N	Y	N	Y	N	Y	N	Y

Notes. Unit of observation is worker in all columns but in Col. 7 and Col.8, where it is discussion group. *Self-reported Engagement* and *Observed Group Activity* are index variables constructed following the methodology from Anderson (2008). The dependent variables in col. 1, col.2 and col.3, col.4 represent the deviation from the factory average of baseline leaders' preferences and views respectively. Probability weights and standard errors clustered at the group level. Controlling for group size FE and stratification FEs (Factory FEs x Union FEs). Selected controls for Col.6 are *Grit* and *Agreeableness squared (BFI)*. *Deviation from Union Preference* is selected for Col.2. No controls are selected for Col.4 and Col.8. R-squared for columns that applied PDS lasso selected controls are estimated by the correlation between the observed outcome and the predicted outcome.

2.14 Appendix: Field implementation

2.14.1 Protocol for random sampling of workers who were not union leaders

We used a random sampling protocol that we designed to obtain a sample that was representative of the target population: sewing operators in the targeted factories, including union members and non-union members. It entailed three stages. First, the CTUM convened the presidents and secretaries of the 28 garment basic unions for an introduction meeting. During the meeting, the CTUM explained the research, requested the unions' participation, and introduced the survey team. Union leaders also completed (1) a factory information form about the factory's sewing lines, their sizes, and their union membership rate and (2) a union information form about the union's organizational structure. Leaders were informed in advance that the survey team would request this information.

Second, the research team matched LLs and EC members to sewing lines and stratified sewing lines by their quartile in the distribution across lines of the share of workers on the line unionized. We then implemented a stratified random selection of up to 11 sewing lines; in factories with fewer than 11 LLs and EC members, the research team selected a number of lines equal to the total number of LLs and EC members. We prioritized LLs, only selecting EC members in factories with fewer than 11 LLs. In factories with fewer than 11 sewing lines, we selected the minimum of {Number of sewing lines, Number of LLs + EC members}. In factories with greater than 50 workers per line, we randomly selected the front or back half of the line to participate. When factories were >80% unionized (<20% unionized), we slightly oversampled lines from the bottom (top) quartile unionization rate. This was to ensure adequate representation of non-union (union) members in field activities. We excluded sewing lines if the president was the only union leader on the line, although, in practice, this was rare.

Third, for each randomly selected line, if it had a LL on it, we assigned the LL to make a complete list of workers on the line, including their union membership status and skill level (higher/low). If a line had multiple LLs, we randomly selected one to make the list. If a line had no LLs, we selected the LL from the nearest non-randomly selected line and broke ties using random selection. We also invited these LLs to participate in the field activities.

LLs brought the lists of workers to their union's first session, which we describe in Section 2.3.2 of the paper. At this stage, the survey team conducted a stratified random selection of around 90 workers per factory; within factory, we stratified by line, union membership, and skill level.

2.14.2 Consensus-building experiment: Discussion prompt provided to groups

At the beginning of the consensus-building experiment, after discussion groups were seated together, the field team explained the prompt below, which they also provided to discussion groups in writing.

> *We are now starting the discussion about minimum wage. Please turn off your phones. The last time the government set the minimum wage was in March 2018. At that time, the government set it at K4800 for an eight-hour workday. The government will announce a new minimum wage in 2020. The CTUM will prepare a proposal for the government on the minimum wage increase. The CTUM wants to gather workers' expectations and opinions to help determine its proposal. For 30 minutes, we would like for you to please discuss the following questions:*
> *(i) How do you think that a minimum wage increase may benefit workers? How do you think that a minimum wage increase may harm workers? Do you think it will affect different groups of workers, for example, skilled versus unskilled, union members versus non-members, differently?*
> *(ii) In 2020, at what level do you think the government will set the new minimum wage for an eight-hour workday?*
> *(iii) In your opinion, what would be the ideal minimum wage level for an eight-hour workday?*
> *Your summary will be provided to the CTUM to help it prepare its proposal to the government. We provide some white blank papers so that you can take notes on these papers while you discuss. At the end of the 30 minutes, please take five minutes to summarize the group's opinions about these questions using this sheet.*

2.14.3 Variable lists

Consensus-building experiment: Engagement Index

- At the end of the discussion, to what extent did your group agree on the prediction of the level of the minimum wage that the government will set?;
- At the end of the discussion, to what extent did your group agree on the ideal level of the minimum wage that the government should set?;
- During the group discussion, I felt confident to express my views and opinions;
- The group discussion was interesting, engaging and informative;
- There were some moments during the discussion when I felt unease and I did not know what to say or do (reversed score);
- All members of my group actively participated in the discussion;
- The group discussion was a waste of my time (reversed score);
- Overall, I enjoyed being part of this group discussion.

Consensus-building: Active Group Index

- Share of workers seem to be engaged in the group discussion (e.g. telling opinions, listening to other people's opinions, writing down notes);
- Share of workers seem to be distracted or not paying attention to the group discussion (e.g. looking down, chatting about irrelevant topics);
- Indicator for one or more persons who are actively facilitating discussion
- Indicator for one or more persons who are asking other workers' opinions
- Indicator for one or more persons who are summarizing group's opinions
- Indicator for one or more persons who are writing down notes

2.14.4 Mobilization Session 3: information provided to workers in each treatment arm

Prior to the surprise invitation, the field team handed the worker their payment in an envelope. After handing them their payment, they read the following scripts:

1. *Leader or staff invitation, no information arm:* Invites worker to do the final survey that is about living standards and working conditions and tells the worker that participation in the survey is entirely voluntary and that it was already very good that they came to the session and did the surveys in the morning. Given that the final survey is a surprise, the research team is going to donate 8000 kyats to buy sewing machines and training fabric for CTUM Training Centre per each discussion group where every member of the group participates in the Minimum Wage Survey.

2. *High coordination information (leader and staff invitation):* Same as (1), plus staff tells worker: "Everyone will be told about the final survey, but LLs might not have time to speak with every worker. They will be able to speak with only X worker in your group," where **X**=group size – 1.

3. *Low coordination information, staff invitation:* Same as (1), plus staff tells worker: "Everyone will be told about the final survey, but LLs might not have time to speak with every worker. They will be able to speak with only **one worker** in your group."

4. *Low coordination information, leader invitation:* Same as (1), plus staff tells worker: "Everyone will be told about the final survey, but LLs might not have time to speak with every worker. They will be able to speak with only **you** in your group."

5. *Social pressure information:* Same as (1), plus staff tells worker: "If you are staying for the survey, I will accompany to the room, and some LLs will welcome you and register you."

Chapter 3

Gender and the Misallocation of Labor Across Countries

This chapter is jointly co-authored with Nava Ashraf, Oriana Bandiera and Victor Quintas-Martinez.

"One of the surest ways to increase national income is to create (. . .) employment for women outside the home."

— Lewis, W.A., 1954. *Economic development with unlimited supplies of labour.*

3.1 Introduction

The gender division of labor inside and outside the home varies across countries and time, but it is always within the confines of norms that assign the largest share of housework to women (Jayachandran (2015); Fernández et al. (2021); Bursztyn et al. (2023)). This under-representation of women in spheres of influence and employment has led to significant effort in both the private and public sectors to address the gap through extensive diversity initiatives (Bertrand, 2020). Critics of these initiatives argue that this can encourage a lower quality bar for minority candidates and, in the corporate world, be ultimately worse for business by hiring or promoting less able candidates. Supporters argue that diversity per se could be beneficial for productivity and profits due, for example, to the nature of the production function or role model effects (Lazear (1999); Athey et al. (2000); Hong and Page (2001)).

If we take as given that the distribution of innate talent is orthogonal to group identity, we cannot understand the effect of diversity on productivity without understanding how under-represented groups select into the labor force. On the one hand, by definition, members of under-represented groups who seek employment likely had to overcome discriminatory barriers, which implies they would be positively selected (Lazear, 2021). This implies that hiring from under-represented groups increases average productivity directly. However, ability is not the only factor that affects selection: differences in other factors unrelated to individual ability- for example, family wealth- could be the reason why individuals are able to overcome barriers and join the labor force.

In this paper, we examine the relationship between gender diversity and performance in the workplace through the lens of selection. To do so, we combine macro-level variation in labor force participation of men and women across cohorts and countries, with micro-level variation in earnings and career paths of approximately 100,000 employees at a large multinational company operating in 101 countries. We are thus able to compare the earnings and the career progression of male and female employees who hold the same job and have similar characteristics such as age, and tenure with the firm, but who faced different barriers when entering the labor market due to their country and cohort of birth.

Our empirical strategy has two prongs. First, we present correlations between women's performance in the firm, measured as wages, promotions, rank, and the female labor force share when their cohort joined the labor market. Under the assumption, underscored by the firm, that HR policies are set centrally and implemented equally in all countries, differences in wages and promotions can be directly mapped to differences in productivity. In other words, if the link between productivity and pay is the same across all countries, we can interpret a negative correlation between performance and female labor force share as a sign of positive selection.

Second, we structurally estimate a simple Roy model of occupational choice to back out individual ability and quantify the productivity cost of the gendered allocation of labor. The structural estimates make use of the fact that we observe several employees in each country, gender, tenure, and cohort, to estimate both a fixed parameter common to all employees in the same cell (e.g. discrimination based on gender) and, using the variation in wages within cell, differences in individual productivity.

To proxy for the barriers that current employees faced when deciding whether to work outside the home, we use the ratio of women to men in the labor force (henceforth LFPR) in the decade when the choice was made. Since we observe employees of different ages in the firm, we can exploit both cross-country and cross-cohort variation in barriers.

The reduced form analysis yields three stylized facts. First, the variation in the LFPR is correlated with the female share of employees in the firm. In other words, in countries and cohorts with low LFPR, the firm hires fewer women. This is in line with the firm using the same selection process in all countries; that is, they do not employ more pro-women policies in countries where norms keep women inside the home (or vice versa). This is also essential for our purposes because if there were no correlation, the variation in LFPR would be moot.

Second, in low LFPR countries and cohorts, women are over-represented in the highest rungs of the hierarchy and are more likely to be promoted relative to their counterparts in high LFPR countries and cohorts. Relatedly, women are over-represented in the top deciles of the wage distribution, and under-represented in the bottom decile, when LFPR is low. This suggests that the women who we observe in the firm in low LFPR countries had the ability to overcome higher barriers, and are thus positively selected.

Third, the wages of the women at the bottom decile of the wage distribution decrease as the LFPR increases, while the wages of the women at the top decile remain constant. This suggests that there is an ability threshold below which women work inside the home and this threshold decreases as LFPR increases. This pattern rules out that the high productivity of women in low LFPR countries, and hence their high wages, are due to the fact that women bring different inputs to the firm and therefore the marginal value of these inputs is high when the share of women is low. If this were the case, we would find that the productivity of the top percentiles of women would decrease as LFPR increased.

A direct implication of these findings is that the gap between female and male earnings is decreasing in LFPR, both because the ability of the average female employee decreases and because the ability of the average male employee increases (Olivetti and Petrongolo, 2008). Since we observe women and men with the same experience, same tenure, and working in the same function, the earnings gap is not influenced by differences in occupational choices that make comparisons between genders difficult (e.g., Blau (1977); Goldin (2014); Card et al. (2015); Wiswall and Zafar (2017); Andrew et al. (2021)).

The findings above are consistent with selection on ability under the assumption that the firm's personnel policy is not correlated with the LFPR. If it were, stronger pro-women affirmative action in low LFPR settings could produce observationally equivalent patterns.

We model the choice between working inside and outside the home in a 2-sector Roy model. Underpinning the 2-sector model is the assumption, supported by empirical evidence, that multinationals pay higher wages and offer more amenities (Hjort et al., 2020) so that all individuals weakly prefer to work for a multinational compared to their next best alternative, be it a local firm, self-employment, or the public sector. We assume that men and women have the same productivity in both sectors and allow for gender-specific barriers to working outside the home.

The value added of the model is that we can separate the effect of confounders common to those of the same gender, such as discrimination, from individual differences in ability.

Individual-level data separately identify gender differences in fixed pay from differences in pay due to differences in productivity. In other words, if we use variation at the cohort-gender-country level we cannot tell apart alternative explanations at the cohort-gender-country level. Our data, however, contains individual-level personnel records and the variation within gender-cohort-country cell allows us to estimate both a cohort-gender-country specific parameter and returns to ability; we are thus able to back out individual ability estimates. Most importantly, this allows us to quantify the productivity cost of barriers faced by women in different countries and to study policy counterfactuals.

We assume that pay is equal to an average wage plus a reward proportional to marginal productivity. This gives us a relationship that links salary in the MNE to the cost of working outside the home. Both variables have a component that is common to all people of the same gender in the same cohort-country-tenure group (for instance, country-specific social norms about gender), and a component that is specific to the individual (their own productivity and preferences). Because we have individual-level pay data, we can separate the common component of salary from the reward for individual-specific performance. We use this to identify productivity thresholds that split the population into workers and homemakers.

We find a gender productivity gap that is positive and large (.9 SD) at low levels of the LFPR and decreases to .3 SD at high levels. The estimated productivity gap is larger in countries with weaker gender equity labor laws and with more conservative gender norms.

Alternatively, we can relax the assumption of equal preferences for work in the two sectors, and use the model to back out the difference in preferences that would make the current LFPR gap optimal. The implied preference gap is several orders of magnitude larger than any other gender gap in preferences estimated in the literature to date.

To test the external validity of our estimates we extract balance sheet data from ORBIS to cover all manufacturing firms in the same countries where the MNE operates: this yields a sample of 2 million firms in 158 SIC3 sectors. We find that the estimate of misallocation we obtain from the MNE data is strongly correlated with the productivity of other firms in the economy, especially those that operate in the same sector as our MNE.

With these estimates in hand, we analyze three counterfactuals. First, given the productivity differences between men and women, we ask whether the firm could effectively "undo" the gender norm by changing the terms of the wage contract to attract more women. To do so, we derive the contract that maximizes productivity while keeping employment and the wage bill constant. We find that the optimal contract has a lower base pay and a steeper performance gradient than the observed contract. This brings the firm's gender ratio close to one and increases productivity by 22%. However, we note that such a contract would massively increase inequality within and between genders; most notably the difference in pay between women and men would go up by 78%. This captures both differences in performance for the same job and differences in jobs as more able women climb the corporate ladder faster. Whilst it is, thus, theoretically possible for the firm to adopt policies that compensate for societal norms, such a steep performance gradient would create a high level of inequality among employees; the highest inequality within the firm would occur where there are the most restrictive gender norms. It may quite possibly also be unsustainable for the firm: in order to hire more women without excessively increasing inequality, they would have to increase women's pay without decreasing men's pay.

The second counterfactual eliminates gender norms: we simulate what would happen if we could eliminate gender differences in the cost of working outside the home. Overall, eliminating gender differences would, by definition, bring the pay gap to zero and would increase productivity by 32% while keeping the wage bill and employment constant. The productivity gains are a result of both high-productivity women joining and low-productivity men leaving. We note that the mirror image of the gender tax that women have to pay to work outside the home is the tax that men have to pay to work inside the home, independent of their skills and preferences. Thus, eliminating gender norms will also eliminate misallocation in work inside the home, by allowing the men who wish to do so to specialize in home production.

The third counterfactual simulates the effect of more stringent labor regulations that make it harder to link pay to performance. We show that this leads to a larger intake of lower-ability workers (Propper and van Reenen, 2010) who, by selection, are more likely to be male in places with more restrictive gender norms.

Our findings show that selection creates a link between the size of a group and the productivity of its members, thus connecting the literature on the barriers to female labor force participation (Goldin (1995); Fernández et al. (2004); Jayachandran (2015); Olivetti and Petrongolo (2016)) to the literature on the impact of diversity for firm productivity (Alesina and La Ferrara (2005); Hamilton et al. (2012); Hjort (2014); Bertrand and Duflo (2017); Marx et al. (2021)). Seen through the lens of selection, the link between diversity and productivity is underpinned by the traits of the minority due to the barriers they had to overcome rather than a direct "treatment" effect of diversity on productivity through, for instance, role model effects or changes in culture. Via selection, the productivity of the firm increases when there are more women in the applicant pool because the firm does not need to hire from the left tail of the distribution of men; a similar pattern is seen in the selection of political candidates in Sweden following the introduction of a gender quota (Besley et al., 2017).

The paper is organized as follows. Section 2 presents the institutional context of the MNE, and describes the data sources. Section 3 introduces the model and provides reduced form evidence on the link between the LFPR and the gender pay gap. In Section 4, we calibrate the parameters of the model from the firm's personnel data and country-cohort level LFPR and Section 5 uses our estimates to evaluate the effect of different counterfactuals. Section 6 concludes by discussing welfare implications and other issues for further research.

3.2 Context and data

3.2.1 Context

We collaborate with an MNE with headquarters in Europe and offices in more than 100 countries worldwide as illustrated in Figure 3.14. The MNE produces consumer goods, in 2019 it had a turnover of €20+ billion and employed over 120,000 workers, of which approximately 55% were white collars. We focus on white-collar workers because blue-collar workers are only observed in two-thirds of countries where the MNE has production activities. Typical white-collar jobs in the MNE involve sales, engineering, marketing, HR, R&D for product development, and general managerial activities. The workers have homogeneous levels of human capital as applications require a college degree, and most employees have degrees in either business administration (50%) or engineering (20%).

3.2.2 Data

Personnel records: Our sample covers the universe of employees between 2015 and 2019. We focus our analysis on local employees (non-expats), resulting in 100,819 distinct regular full-time workers over 2015-2019 in 101 countries (303,756 employee-year observations). The company is organized into a hierarchy of work levels that goes from work level 1 to 6 (C-Suite). For each employee, we observe: (1) their work level, (2) their tenure in the firm and post, and (3) total compensation (fixed plus variable pay in euros). Pay captures differences in performance between employees, encompassing off-peak salary increases as well as promotions. We look at four 10-year age cohorts within the company, 18-29, 30-39, 40-49, and 50-59[1]. Thanks A recent strand of evidence shows that multinationals pay higher wages in developing countries (Hjort et al. (2020); Alfaro-Urena et al. (2019)), and Appendix Figure 3.15 shows that the firm's average wages are usually well above the countries' average wages using both the average wages in the manufacturing sector from the ORBIS database and the ILO estimates for white-collar employees. Table 3.1 presents summary statistics separately by gender at the gender-cohort-

[1] We do not have more granular data on age cohorts because of data privacy clauses. Due to a limited sample size of workers in age cohorts above 50-59, we only consider workers up to the 50-59 cohort.

country-tenure level, which is the relevant unit of analysis used in the structural estimation.[2] The female cells show lower average pay, age, tenure, the probability to be in managerial positions, and the probability to experience fast promotions. Overall in the company, 40% of workers are aged between 30-39 and the majority of workers are in WL1 (>70%).

Country-cohort level data: We combine the firm's administrative records with country-cohort data on labor force participation rates of men and women from the World Bank. In particular, we match the age cohorts in the firm with the average LFP rate in the country in the decade of labor market entry, separately by gender. For example, employees of age 18-29 are associated with the LFP rates of the 2010-2020 decade while employees of age 30-39 are associated with the LFP rates of the 2000-2010 decade, and so on.[3]

Figure 3.1 shows moments of the distribution of the labor force participation ratio of women's LFP against men's LFP (LFPR)[4] at different deciles of GDP per capita and shows that while the mean exhibits the well-known U-shape pattern (Goldin, 1995), the distributions largely overlap: there is variation in LFPR at every level of GDP per capita. For instance, the interquartile range of LFPR is broadly similar across the deciles of GDP per capita. Moreover, the 90th percentile of LFPR only ranges between 1 and 0.9. This indicates that there are country-cohort cells with high LFPR at every level of economic development, not solely concentrated in high-income countries. This is essential for the analysis that follows because it allows us to partial out the differences in national income among the countries where the MNE operates.

Figure 3.2 documents a large variation in female shares of employment within the MNE across countries, which closely matches the variation in LFPR across countries. The relationship between the firm and country ratios does not vary with the level of LFPR, which is confirmed by a formal test of differential slopes by above/below median LFPR. This suggests that the selection of men and women in the MNE follows closely their LFP decisions and hence that the countries' LFPRs "bind".

[2] As we explain in section 3.4, we let the structural parameters vary by gender, cohort, country, and tenure in the firm to control for confounders and allow for the wage policy to differ across countries as well as take into account factors such as on-the-job learning.

[3] Since data from the World Bank on the LFP for the 1980-1989 decade is not available, we use a linear interpolation of the LFP for that decade.

[4] We rescale women's LFP my men's LFP in order to account for country-cohort level differences in the probability of individuals to work. However, we note that the variation in men's LFP is tiny compared to the variation in women's LFP. In other words, the variation in LFPR is mostly driven by the variation in the female labor force participation, the numerator in the ratio.

3.3 Diversity and productivity through the lens of selection

3.3.1 Framework

To formalize our intuition about differential selection into the labor force, consider a basic two-sector Roy (1951) model (as formalized by Borjas (1987)). Suppose that the utility from working outside the home is equal to pay and that these are a (log-)linear function of individual i's productivity, A_i:[5]

$$y_i^1 = \alpha^1 + \beta^1 A_i, \tag{3.1}$$

The term α^1 is the unconditional average wage and β^1 is the return to productivity. We interpret deviations from that average wage as arising from individual differences in productivity.

Similarly, we model worker i's value of housework as:

$$y_i^0 = \alpha^0 + \nu^0 N_i, \tag{3.2}$$

where N_i captures sources of individual heterogeneity in the value of staying out of the labor force that are, without loss of generality, independent of productivity. Here, α^0 captures the unconditional average value of staying out of the labor force (e.g. social norms that affect all women).

We make the distributional assumption:

$$\begin{bmatrix} A_i \\ N_i \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right)$$

equally distributed in all our sample. A difference with the canonical Roy-Borjas model is that, as a result, y_i^1 and y_i^0 are independent. This assumption will allow us, in section 3.4, to identify the parameters of the model with individual-level data only on those who are working at the firm. It has also become a standard assumption in the literature on misallocation with structural models of sectoral choice (e.g. Hsieh et al. (2019)).

Individual i self-selects to work outside the home if and only if:

$$y_i^1 \geq y_i^0$$

$$\Leftrightarrow \quad \eta_i \equiv \frac{\beta^1 A_i - \nu^0 N_i}{\sqrt{(\beta^1)^2 + (\nu^0)^2}} \geq \frac{\alpha^0 - \alpha^1}{\sqrt{(\beta^1)^2 + (\nu^0)^2}} \equiv \xi.$$

[5] This could be micro-founded, for example, by assuming that workers are paid their marginal product of labor, that all workers supply the same amount of hours (full-time) and that the individual production function is Cobb-Douglas, $F(K, l_i) = e^{z_i} K^\alpha l_i^{1-\alpha}$, where $z_i \sim \mathcal{N}(\mu, \sigma^2)$.

Because $\eta_i \sim \mathcal{N}(0,1)$, that happens with probability $1-\Phi(\xi)$. It is straightforward to show that:

Result 1: *Stronger norms about gender roles, that is a higher α_0, increase ξ and therefore reduce the probability of working outside the home $1-\Phi(\xi)$.*
Moreover, we have that:[6]

$$\begin{aligned}\mathbb{E}[A_i \mid \eta_i \geq \xi] =& \frac{\operatorname{cov}(A_i,\eta_i)}{\operatorname{var}(\eta_i)}\mathbb{E}[\eta_i \mid \eta_i \geq \xi]\\ =& \frac{\beta^1}{\sqrt{(\beta^1)^2+(\nu^0)^2}}\frac{\phi(\xi)}{1-\Phi(\xi)}.\end{aligned}$$

This implies that:

Result 2: *The average ability of those who choose to work outside the home is increasing in ξ.* Taken together with the previous result this underpins the link between productivity and diversity: higher barriers lead to fewer individuals joining the labor force and their average ability to be higher, all else equal.

3.3.2 Female share and performance

Figure 3.2 shows that the firm hires more women as the LFPR increases: the female to male ratio in the firm closely follows the same ratio in the labor force. Under our framework, an increase in the proportion of women in the firm, keeping firm size constant, would imply that the marginal woman hired is less productive than the average woman while the man she replaces is less productive than the average man. Hence, we would expect the average performance of women to fall and the average performance of men to rise as the LFPR increases.

Figure 3.3 displays plots of women's performance measures compared to men's against the labor force participation ratio. In line with this intuition, Panels (a) and (b) of Figure 3.3 show that women are over-represented at the top decile of the wage distribution and under-represented at the bottom decile when overall LFPR is low and converge when LFPR is close to 1. Moreover, Panels (c) and (d) show that when LFPR is low women are over-represented in managerial positions and among those who get promoted quickly, but the two converge as participation rates get more equal. We only show the plots with the women share as the y-axis, since by definition, the plots with the men share would be their symmetric equivalents.

Our framework indicates that as the LFPR increases, women with lower ability enter the

[6] To see why, realize that we can write $A_i = \operatorname{cov}(A_i,\eta_i)/\operatorname{var}(\eta_i) \times \eta_i + u_i$, where u_i is the population OLS projection error. By the normality assumption, u_i and η_i are not just uncorrelated, but also stochastically independent.

labor force while high ability women already in the labor force are unaffected. We test this result empirically in Figure 3.4 by looking at how women's wages at different points of the distribution change as the LFPR increases. Since overall wage levels change across countries, we control for the respective wage measure for men. The first panel from the left shows that women's average wages decrease as the LFPR increases. The remaining two panels indicate whether this decrease is driven by the bottom or the top of women's wage distribution. It comes from the bottom of the wage distribution: the 10th percentile of women's wages decreases with LFPR (second panel), while there is no impact on the 90th percentile (third panel).[7]

These differential patterns at different levels of the wage distribution are consistent with women facing a higher bar for participating in the labor force in countries with lower LFPR. As the LFPR increases, lower ability women start to enter, hence decreasing the wages at the bottom of the distribution, while high ability women are not affected, hence leaving the wages at the top of distribution unchanged.

Figure 3.4 is also useful to rule out an alternative interpretation that posits that women's numerical scarcity leads to women earning wage premia in environments where women are scant. According to this interpretation, women receive high wages in low LFPR settings not because employed women are more able in general, but rather because female-specific skills are scarce and hence command a higher price. Specifically, women's contribution would be particularly valuable in a low LFPR environment, where most employees are men. If this were true, women's wages should be negatively correlated with the LFPR, and even more so for the wages of high ability women compared to low ability women, which is not what we find in Figure 3.4.

3.3.3 Implications for the gender pay gap

In light of the selection model in subsection 3.3.1 and the empirical evidence in the previous subsection 3.3.2, we infer that barriers to entry are the dominating force in generating gender differences in LFP. In countries with low LFPR the productivity threshold for women to work is so high that only the most talented women exceed it. As a result, the lower the LFPR, the more positively selected the women hired within the MNE are, compared to men.

The differences in the wage data of female and male workers as the LFPR decreases have implications on how the gender pay gap changes with the LFPR. Assuming the gap is near zero when the LFPR is 1, it should become larger (more positive) as the share of women in the labor force decreases.[8] Namely, the correlation between the gender pay gap and the LFPR should have a negative sign.

[7] Results are robust to using other percentiles, for example the 25th percentile.

[8] We define the gender pay gap as the ratio of women's wages over men's wages.

Table 3.2 combines women and men's wages together and shows that the correlation between the gender pay gap and the LFPR is indeed negative. It shows the gender pay gap using different sources of variation, thus also allowing us to check whether the sign of the correlation is consistent regardless of the source of variation that we use to identify it. We estimate the following model:

$$w_{iact} = \alpha LFPR_{ac} + \beta Female_i + \gamma LFPR_{ac} * Female_i + \mathbf{X'}_{iact}\boldsymbol{\Lambda} + \boldsymbol{\psi}_t + \epsilon_{iact} \tag{3.3}$$

where w is log wage of employee i in country c and year t for age group a. $\boldsymbol{\psi}_t$ represents year fixed effects to take out year-level macro shocks and $\mathbf{X}_{iact}$ is a vector of controls. We cluster standard errors at the same level as the RHS variable, that is country-cohort. The coefficient of interest is γ which measures the change in the pay gap as LFPR increases and $\gamma < 0$ would indicate the presence of differential barriers to entry for women.

We include different controls in $\mathbf{X}_{iact}$: column 2 controls for a quadratic function of tenure and function fixed effects, column 3 adds GDP per capita in logs so to show that the variation in the LFPR is not only a function of country income, column 4 adds cohort fixed effects so to only exploit the variation across countries; column 5 replaces the cohort fixed effects with country fixed effects hence only exploiting the variation within countries. The comparison between columns 4 and 5 is particularly informative as it uses one source of variation at a time.

The estimates of γ are negative and precisely estimated in all specifications. In Appendix Table 3.6, we use the LFP data for individuals with advanced education only.[9] The results are nearly unchanged when we adopt this measure. However, we lose almost 20% of the sample and particularly from countries with low female labor force participation (FLFP). Hence, we employ the overall LFP as our default measure.[10] The fact that the γ coefficient is stable does reassure us that it is, in fact, measuring selection.[11]

Finally, columns 6 and 7 estimate the same specification as in column 1 to look at the pay progression for new hires. They show that women in low LFPR country-cohorts display faster pay growth and a higher probability of promotion and that this positive gender gap in

[9] As defined by the World Bank, an individual with advanced education has completed a short-cycle tertiary education or a college degree and/or above.

[10] The correlation between overall LFPR and LFPR for individuals with advanced education is 68%.

[11] We conduct a number of additional robustness checks. Table 3.7 in the Appendix shows that the patterns in Table 3.2 hold if we add fixed effects for the geographical region and if we split the sample by lower income and higher income countries (as defined by the World Bank). In Appendix Table 3.8 we report the results when converting wages from euros into PPP 2017 \$, using the PPP conversion rates of the ICP at the World Bank. The gender gap is unaffected and the only change is the magnitude of the coefficient on LFPR which shrinks to the level found when controlling for country fixed effects (column 5 in Table 3.2). This is what we would expect as differences in PPP exchange rates would not affect cross-country comparisons of the gender gap. Finally, we note that the results are driven by differences in fixed pay rather than in variable bonus, which constitutes a much lower proportion of overall salary (the median ratio of bonus to fixed pay is 13%) — see Appendix Table 3.9. In the company, pay summarizes altogether most differences in performance between employees, encompassing off-peak salary increases as well as promotions.

realized performance decreases as the LFPR increases. This evidence on the sample of new hires also supports the interpretation that the firm is not taking into account the implications on productivity that different LFP rates by gender might entail when making hiring and wage offers decisions. It is only after some time at the firm that women and men *ex-post* pay gaps start to diverge the lower the LFPR is, indicating positive selection of women into the firm due to ex-ante unobservable characteristics.

3.4 Quantifying misallocation

In this section, we take the model of worker selection into the labor force that we introduced in section 3.3 to the data. The value added of the model, relative to the reduced form estimates and to the existing literature, is that we can leverage our individual-level data to separately identify gender differences in fixed pay (due to -for instance - discrimination) from differences due to productivity. Normally this cannot be estimated with aggregate data as it only contains the average wage while individual data can be used to compute the variance which maps directly to the variation in pay due to differences in productivity. We first discuss how to calibrate the parameters of the model from the firm's personnel data and country-cohort level LFP. Next, we present the results of our calibration and validate them against different variables and independent datasets.

3.4.1 Model calibration

For our structural exercise, we will use the two-sector Roy-Borjas model presented in section 3.3. We let the parameters vary by country (c), cohort (a), and tenure in the firm (t)[12] and gender (g) cells: $\alpha^1_{gtac}, \beta^1_{gtac}, \alpha^0_{gtac}, \nu^0_{gtac}$. This flexibility allows us to control for a variety of confounders: for example, we do not make assumptions as to whether the firm has the same wage policy for men and women or across countries, and the fact that we estimate β^1 by tenure group allows for on-the-job learning in a way that is possibly correlated with productivity. We keep the normalization that (A_i, N_i) have the same jointly standard normal underlying distribution across *gtac* cells since differences in that distribution would be absorbed in our calibration by $\alpha^1_{gtac}, \beta^1_{gtac}, \alpha^0_{gtac}$ and ν^0_{gtac}.

In what follows, we make three assumptions. First, as with any other binary choice models, the scale parameter of the selection Probit of our Roy model is not identified. Hence, we normalize it to 1 for calibration purposes, so that $\mathrm{Var}(\eta_{igtac}) = \mathrm{Var}(\beta^1_{gtac}A_i - \nu^0_{gtac}N_i) = (\beta^1_{gtac})^2 + (\nu^0_{gtac})^2 = 1$.

[12] We aggregate tenure in groups of two years each.

Second, we do not have individual-level data for those not working at the firm, so we are not able to identify the correlation between A_i and η_{igtac} (as one would do with a standard Heckman (1976), selection model). Hence, we are going to assume that A_i and N_i are independent.

Finally, again because we have data only on those in our firm, we need to obtain a measure of the probability that a worker will want to work for our firm, i.e. $\Pr(y^1_{igtac} \geq y^0_{igtac})$. We proxy that by country-cohort-level LFP data. In other words, we are identifying the probability that an individual chooses to work for our firm with the probability that he or she chooses to work for *any* other firm. For that, we need to assume that working at our firm is weakly preferred to working at other firms for all individuals. We provide supporting evidence of this in section 3.3.

Under our distributional assumptions, labor force participation and the observed wage satisfy the following moment conditions:

$$\Pr(\text{employed}) = 1 - \Phi(\xi_{gac}) \tag{3.4}$$

$$\mathbb{E}[y^1_{igtac} \,|\, \text{employed}] = \alpha^1_{gtac} + (\beta^1_{gtac})^2 \lambda(\xi_{gac}) \tag{3.5}$$

$$\text{Var}(y^1_{igtac} \,|\, \text{employed}) = (\beta^1_{gtac})^2 + (\beta^1_{gtac})^4 \left[\xi_{gac}\lambda(\xi_{gac}) - \lambda(\xi_{gac})^2\right] \tag{3.6}$$

where $\lambda(\cdot) \equiv \phi(\cdot)/(1-\Phi(\cdot))$ is the inverse Mills ratio and $\xi_{gac} \equiv \alpha^0_{gtac} - \alpha^1_{gtac}$ are the participation thresholds. Together with the restriction $(\beta^1_{gtac})^2 + (\nu^0_{gtac})^2 = 1$, we have four equations in four unknown parameters $(\alpha^1_{gtac}, \beta^1_{gtac}, \alpha^0_{gtac}, \nu^0_{gtac})$ for each gender, tenure level, age cohort, and country cell.

In order to calibrate those parameters, we match the moments above to their empirical counterparts. Since we observe the wage for those working in the firm, we can use the sample average and variance as the empirical analogs for $\mathbb{E}[y^1_{igtac} \,|\, \text{employed}]$ and $\text{Var}(y^1_{igtac} \,|\, \text{employed})$. To eliminate the effect of observables, we use the residuals of a regression of log(base pay + bonus) on year and function dummies as our measure of y^1_{igtac}. To calibrate the parameters in the participation decision, we use World Bank LFP data in each gender, cohort, and country cell, LFP_{gac}, as our empirical analog for Pr(employed). Table 3.3 provides a summary of the parameters of the model and the empirical target that each parameter tries to match in our calibration strategy.

3.4.2 Calibrated parameters

Once we have obtained the parameters in the firm's wage policy, we can recover productivity as

$$\widehat{A}_i = \frac{y^1_{igtac} - \alpha^1_{gtac}}{\beta^1_{gtac}},$$

where y^1_{igtac} is the residualized log-wage described in section 3.3. Figure 3.5 plots the average calibrated productivity for our sample of firm workers by LFPR. Average productivity is approximately constant for male workers in our countries, whereas for female workers, average productivity is very high when FLFP is much smaller than MLFP (about 0.8 standard deviations higher than for men) and decreases as the LFPR approaches 1. This result is consistent with selection with respect to productivity, so the lower a group's LFP, the more positively selected they are.

Figure 3.6 further shows that this is due to a shift in the whole productivity distribution. For men, the distribution of calibrated productivity is very concentrated and somewhat left-skewed for the lowest levels of LFPR, becoming more dispersed as the LFPR increases. For women, as the LFPR increases, there is a downward shift of the entire distribution but the right tail.

If we believe that the underlying distributions of A_i and N_i in the population are the same for men and women, differences in LFP across genders must be due to the cost of social norms. In other words, if men and women have the same distribution of productivity[13] and preferences for staying at home, the only reason for their LFP to differ is that they must face different payoffs from working outside the home.

3.4.3 Productivity or preferences?

Alternatively, we can relax the assumption of equal preferences for housework and use the model to back out the difference in preferences that would make the current LFP gap optimal. If the observed LFPs are optimal so that the marginal man and woman have the same productivity, we should have $\mu^F_N - \mu^M_N = \Phi^{-1}(MLFP) - \Phi^{-1}(FLFP)$, where μ^g_N is the average N_i of gender g (assuming that the standard deviation is the same).

Figure 3.7 plots these average preference gaps by country. Two points are of note: first, differences between genders are large - well over $1S.D.$ for the top decile of countries, and at least $0.5S.D.$ for most of the sample. Second, differences across countries are large, for instance, the interquartile and the interdecile ranges are $0.45S.D.$ and $1.18S.D.$, respectively. These patterns are in sharp contrast with the existing evidence on gender differences in economic preferences across countries where they are much smaller. Figure 3.16 in the appendix shows the distribution of gender differences in risk aversion, altruism, trust, patience, positive and negative reciprocity from Falk et al. (2018) for the same set of countries. Overall, considering the interquartile range, for differences in preferences to rationalize our results we would need these to be at least three times larger than any other economic preference.

[13] Gender differences in educational attainment have been drastically reduced over the last decades and are much smaller compared to the LFP gap (see Appendix Figure 3.20).

3.4.4 Validation

To validate our estimates we show that they correlate with other data not used to calibrate the model. We use three sets of external variables to validate our three key estimates: (i) individual performance data from the firm's records to validate calibrated productivity; (ii) country-level labor laws that constrain the firm personnel policy to validate the parameters of the wage policy and (iii) country-level social norms to validate our estimate of the difference in home payoffs.

The results for the first exercise are in Appendix Figure 3.17. Panel (a) validates calibrated productivity against the firm's performance appraisals that a manager gives every year, panel (b) against pay growth in the first year (for new hires), and panel (c) against an objective productivity measure for the sales department (based on reaching specific targets). The correlation of all three indicators with our calibrated productivity is positive and strong.

Appendix Figure 3.18 panel (a) shows the results for the second exercise. It plots our calibrated β^1_{gtac} (which represents returns to productivity, and, in our model, is what generates dispersion in pay within gender-country-cohort-tenure cells) against the Restrictive Labor Regulations Index from the World Economic Forum.[14] Consistent with stricter labor regulation limiting performance pay, we find that our calibrated β^1_{gtac} is lower in countries with a higher value of the index.

Appendix Figure 3.18, panel (b) plots the gap in our calibrated $\alpha^0_{Ftac} - \alpha^0_{Mtac}$ (that we interpret as the cost of gender norms for the average woman) against the Women, Business and the Law Index from the World Bank.[15] The figure shows that the gap $\alpha^0_{Ftac} - \alpha^0_{Mtac}$ is strongly negatively correlated with laws allowing or facilitating women's labor. Therefore, part of the restrictions to FLFP due to gender norms may actually be embedded in the laws of certain countries.

Finally, Appendix Figure 3.19 shows the gap in our calibrated average value of staying at home, $\alpha^0_{Ftac} - \alpha^0_{Mtac}$, against the responses to four questions in the World Value Survey: (1) "Men make better business executives than women do," (2) "Pre-school child suffers with working mother," (3) "Being a housewife is just as fulfilling as working," (4) "When jobs are scarce, men should have more right to a job than women." For all four questions, we see a strong

[14] The WEF Restrictive Labor Regulations Index is available for the period 2008–2020 and it is based on an annual survey on the most problematic factors for doing business (e.g. corruption, taxes, inflation, etc.). The survey is administered to a representative sample of around 15,000 business executives in 150 countries. The Restrictive Labor Regulations Index includes measures related to labor-employer relations, wage flexibility, hiring and firing practices, performance pay, labor taxes, attraction and retention of talent.

[15] The WB Women, Business and the Law Index covers 190 countries through the period 1971–2020 and is structured around the life cycle of a working woman. It consists of eight indicators constructed around women's interactions with the law — mobility, pay, workplace, marriage, parenthood, entrepreneurship, assets, and pensions — for current laws and regulations (i.e. religious and customary laws are not considered unless they are coded). Hyland et al. (2021) provides an overview of the data documenting how gender discrimination by law affects women's economic opportunities.

positive correlation between agreement with the statement and our gap in the average value of staying at home.

3.4.5 Beyond the firm

To gauge the external validity of our estimates we extract balance sheet data for all manufacturing firms from the ORBIS database for our sample countries and years, and we test whether our estimate of the gender productivity gap (the average productivity of women minus the average productivity of men at the country level) correlates with the productivity of other firms in the economy. We obtain a cross-section of firms based on the latest year of balance sheet reporting between 2012 and 2019. The sample contains two million firms in 158 SIC3 sectors. Intuitively, our estimate of the gender productivity gap captures differences in productivity due to differences in barriers at the societal level as well as idiosyncratic differences due to the MNE pay and recruitment policies. If the latter dominates, the estimated productivity gap will not be correlated with the productivity of other firms. Table 3.4 reports the estimates of the following model:

$$y_{isc} = \alpha l_{isc} + \beta k_{isc} + \gamma(A_F - A_M)_c + \delta X_c + \epsilon_{isc} \tag{3.7}$$

where y_{isc}, l_{isc}, k_{isc}, are, respectively, log operating revenue, log employment, and log capital of firm i in sector s in country c. Since our estimated productivity gap varies at the country level, we control for GDP and the LFPR in X_c and cluster standard errors by country. The coefficient of our estimated productivity gap is negative, precisely estimated, and quite large, likely due to correlated unobservables at the country level. To estimate correlations within country, we split sectors into two groups: those where the MNE operates and those where it does not. In column 3 we estimate a fixed effect model that identifies the correlation between our productivity gap and the productivity of firms in the same sector minus the correlation with the productivity of firms in different sectors. The estimated elasticity in column 3 is about half of that of capital.

Table 3.5 estimates the model in Equation 3.7 at the sector level, on both the mean and the variance of the productivity gap. In line with the findings at the firm level, average productivity is negatively correlated with the calibrated productivity gap. Moreover, the dispersion of productivity between firms in the same sector and country (a rough measure of misallocation) is higher in countries where the productivity gap is higher.

More generally our data is not well-suited to evaluate welfare impacts. Eliminating the norms that lead to higher barriers for women has two consequences for welfare. First, the reallocation of labor inside and outside the home creates winners and losers. The former consists of the women who are strongly suited for work outside the home and transition from inside to outside as well as of the men who are strongly suited for work inside the home who move in the opposite direction. However, men who get crowded out by women in the workplace might lose. The extent to which this happens depends on the second consequence of eliminating misallocation, which is the increase in efficiency and productivity overall, which generates a virtuous circle through increases in demand. We cannot assess these effects because we do not know the production function of the firm, so we cannot say how the firm would adjust employment in response to a convergence of gender norms. In summary, we are able to quantify the effect of misallocation on the share of the pie that everybody gets but we cannot quantify the effect of misallocation on the size of the pie.

3.5 Counterfactuals

We use the model estimates to evaluate the effect of different counterfactuals on the LFPR, the pay gap, and welfare. To do this we need to take a stance on how the firm responds to changes in the environment. Since we do not observe the production function of the firm nor the elasticity of demand they face we cannot use profit maximization as the guiding criterion. Rather, we take the observed level of employment and the wage bill in each country-cohort-tenure group as binding constraints.[16] Our first counterfactual asks whether, under these constraints, the firm maximizes productivity. Secondly, we quantify the effect of misallocation on the firm's productivity. In the third counterfactual, we simulate the effect of stricter labor laws.

3.5.1 Does the firm maximize productivity?

To compute the optimal wage policy, we let the firm choose α^1_{gtac}, β^1_{gtac} to maximize the productivity of its employees, subject to two constraints: (i) keeping total employment (or the labor force) fixed, (ii) keeping the total wage bill constant.[17] The exact details of the maximization problems are discussed in Appendix 3.10.3.

Figure 3.8, panels (a) and (b) compare the calibrated wage policy parameters to the solution of the optimization problem described above. We can see that these do not coincide, and for some countries, they are quite far apart. The difference between the optimal and the observed parameters follows the same pattern in most countries: to maximize productivity the firm should increase the fixed pay of women α^1_F, decrease the fixed pay of men α^1_M and increase variable

[16] In practice these are determined by a maximization problem that we do not observe. This is equivalent to assuming that the firm sets the optimal scale of operation and then decides who to hire to maximize productivity.

[17] Both constraints are necessary to make sure that we obtain a sensible solution. Without the employment constraint, the firm can increase average productivity just by hiring fewer people (because of positive selection). Without the wage bill constraint, the firm can raise both α^1_{gtac} and β^1_{gtac} in a way that increases productivity without changing employment, at the cost of paying much higher wages.

pay β^1 for both genders.

Under the optimal policy, the firm equates the participation thresholds of men and women, so that the LFPR gets close to 1 as shown in Figure 3.8 panel (e). While equating marginal productivity across groups is obviously the solution to the unconstrained problem, this exercise tells us that it is possible to do so while simultaneously keeping employment and the wage bill constant. The optimal policy effectively undoes differences in LFP and leads to higher productivity in every country as shown in Figure 3.8 panel (c). On average, the firm could increase productivity by 22%.

Why is the firm not setting the optimal α^1_{gtac}, β^1_{gtac}? A possible answer is that while the average pay is constant by assumption, the optimal policy generates a stark increase in inequality between genders. Indeed, because women in the labor force are more positively selected in most countries, it would be productivity-maximizing to pay women more, both in terms of fixed and variable pay, so that the pay gap between females and males would be even larger than what we observe. Without any change in norms, on average over all countries, the gender pay gap (Female – Male) would have to increase by 78% (Figure 3.8, panel (d)). Since individual productivity is not directly observable, such an increase might not be acceptable.

Another important reason may be labor regulations, which limit variable pay and pay inequality even within gender. We address this in subsection 3.5.3.

Finally, the fact that the firm would be better off hiring more women (in most cases, hiring as many women as men) suggests that quotas would not bind. However, meeting them would require a large increase in inequality between genders, with steep rewards for talent, in order to sufficiently attract women into the labor force. This could be as stark as, for example, most leadership positions being held by women while all men would work as subordinates. Note that this would be a very different scenario than that of many policies which prescribe equality in pay and rewards between genders, and would imply an equal number of men and women in top-level positions, as well as in lower-level positions.[18]

While our data is well suited to quantify the cumulative productivity loss due to the gendered division of labor, it does not shed light on the individual components that make up gendered labor division, such as barriers to education, how work affects marriage prospects, child penalties, and so on. We calibrate gains from eliminating gender norms at every stage; if we eliminated barriers to LFP without eliminating barriers to education, the gains would of course be smaller because the pool of qualified women applicants would be smaller. We thus do not know ex-ante

[18]If the policy were to combine gender quotas with the imposition of equal pay, rather than rewarding talent regardless of gender, the firm will reduce mean productivity to minimize labor costs. Results are available on request.

what the size of the pool of qualified women would be if certain gender barriers were eliminated but others were not. This precludes us from analyzing the trade-off between diversity and quality that is at the core of popular policy tools such as quotas and affirmative action

3.5.2 How much does misallocation cost?

Our next counterfactual quantifies the effect of misallocation on the firm's productivity and on workers' welfare. In our framework the true value of staying at home for a woman equals that of a man with the same observable characteristics and (A_i, N_i) type, $y^{0,*}_{iFtac} = y^0_{iMtac}$. However, when entering the labor force, women must pay a "gender norms tax" τ_{tac} (as in Hsieh et al. (2019)) proportional to their true value of staying at home, so that when making the decision of whether to enter the labor force or not they take into account $y^0_{iFtac} = (1 + \tau_{tac})y^{0,*}_{iFtac}$. In this counterfactual, we eliminate the gender norms tax by setting the value of the staying at home parameters of women (α^0, ν^0) equal to those of men (within the same country, cohort, and tenure cells).

We discuss effects both in the short run, that is, keeping the pay policy of the firm fixed, and in the long run, when the firm can optimally adjust its policy to the new environment. To do this we must take a stance on the firm's optimization procedure. In light of the results in subsection 3.5.1 we let the firm choose α^1_{gtac}, β^1_{gtac} to maximize the productivity of its employees, subject to three constraints: (i) keeping total employment (or labor force) fixed, (ii) keeping the total wage bill constant, (iii) a bound on the pay gap in the firm. This bound is chosen to maximize the goodness of fit between the constrained optimal and calibrated wage policy parameters at baseline. The exact details of the maximization problems are discussed in Appendix 3.10.3.

Figure 3.9 plots the short-run effects of eliminating gender norms on the LFPR, the pay gap, and average productivity. In the short run, MLFP does not change (because men's value of staying at home and wages stay the same), and FLFP increases, so LFP increases overall, which means that the ratio increases as well. These changes in LFP reduce the pay gap because the women that enter the labor force are less able on average than those who were already working at baseline, so women's average wages decrease. For the same reason, average productivity decreases, since MLFP does not change, but FLFP increases, and the entrants are, on average, lower productivity than those who were already in the labor force. These findings mirror those in subsection 3.5.1: the baseline policy sacrifices productivity to bound pay inequality when entry barriers vary by gender, and hence, once barriers are equalized, the productivity cost rises.

The trade-off between productivity and inequality when the gender tax is set to zero however disappears in the long run (Figure 3.10). If the firm can adjust its policy, it will equate the participation thresholds so that the LFPR becomes 1. The fixed total employment constraint means that MLFP decreases by the same amount as FLFP increases. In the long run, because we equate the parameters of staying at home, it is optimal for the firm to have the same wage policy for men and women, which eliminates the pay gap. Again, the constant wage bill constraint means that this is just a redistribution from men to women. Finally, in a world without norms tax, the firm replaces less able men with more able women, and average productivity increases.

Figure 3.11 shows, by baseline LFPR, the average productivity of LFP entrants and leavers. The average productivity of LFP entrants is, because of positive selection, lower in those countries where the LFPR changes are the largest. However, those entrants are on average much more productive than the leavers they replace. Figure 3.12 shows the productivity gains of eliminating gender norms by countries. The average across countries is 32%, although there is substantial heterogeneity, with some countries, such as Pakistan or Sri Lanka, obtaining potentially an increase of up to 90% in productivity.

3.5.3 Stricter labor laws

In a third counterfactual, we ask what the effect of stricter labor laws, that limit performance pay, would be on average productivity in the firm. To answer this question, we consider the constrained optimal wage policy and add an additional cap on returns to productivity: $\beta^1_{gtac} \leq \max\{\beta^1_{Fta,FRA}, \beta^1_{Mta,FRA}\}$, where $\beta^1_{Fta,FRA}, \beta^1_{Mta,FRA}$ denote the corresponding parameters for France (since France is the 95th percentile of the WEF Restrictive Labor Regulations Index). We leave the value of staying-at-home parameters $(\alpha^0_{gtac}, \nu^0_{gtac})$ unchanged at baseline, but we allow for the LFP of each group to respond optimally to the change in incentives.

The results are plotted in Figure 3.13. The average productivity of women does not change, whereas the average productivity of men decreases substantially. Intuitively, limiting performance pay hurts the firm because they cannot screen out lower-productivity men, whereas it does not affect women because lower-productivity women are out of the labor force even at baseline.

3.6 Conclusion

We have documented that the gendered division of labor between the market and the home creates a direct link between diversity and productivity through selection. Evidence from white-collar workers in a large multinational that operates around the world has shown that women are more positively selected in places where female labor force participation is low: the average productivity of female employees falls as FLFP rises. When FLFP is low, the firm has to hire mostly men, leaving significant productivity gains unrealized. Understanding differential selection by gender, or indeed by any under-represented group, is key to informing personnel policy as well as broader labor market policies.

Our findings have three implications for the design of firms' diversity policies at the hiring and compensation stages. The first implication is that even if the firm does not observe potential productivity at the point of hiring, the logic of selection indicates that under-represented groups, other things equal, will likely have higher productivity. This implies that between two potential hires with the exact same observable qualities, the minority candidate will have the better unobservables on average. This leads to the second implication, namely that diversity is justifiable on the grounds of productivity, regardless of whether diversity has an independent effect of its own on productivity. The third implication is that aiming for gender equity — in pay, promotions, and dismissals — can turn out to be inequitable because selection generates different distributions of productivity. Perhaps counter-intuitively, gender equity policies might end up hurting women as they limit the firm's ability to reward performance. As things stand, greater rewards for objective performance could induce more women to enter the labor market (and stay there). Rewards need not be monetary; indeed in the presence of social norms where women may bear a disproportionate burden of childcare and there are transaction costs, resources that would enable women to better manage childcare could have equally large productivity gains for the firm through the selection margin.

Awareness of the evidence of positive selection of under-represented groups changes how quality can be inferred, particularly if quality is not perfectly observable or objectively measured: the very presence of a member of an under-represented group should change one's prior on the talent of that member, simply through the logic of positive selection. This could lead to greater expectations, possibly in the form of greater compensation, placed on minority candidates which could have opposing effects: on the one side, it could put greater pressure on minority candidates to perform or de-motivate non-minority employees and, on the other side, it could provide greater motivation for diverse employees by being received with positive expectations by management. The implications of increased awareness of the positive selection effect of diverse candidates by leaders for the recruitment, promotion, and productivity of under-represented groups are left to future research.

3.7 References

Alesina, Alberto and Eliana La Ferrara, "Ethnic diversity and economic performance," *Journal of economic literature*, 2005, *43* (3), 762–800.

Alfaro-Urena, Alonso, Isabela Manelici, and Jose P Vasquez, "The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-firm Linkages," *Working paper*, 2019.

Andrew, Alison, Oriana Bandiera, Monica Costa-Dias, and Camille Landais, "Women and men at work," *IFS Deaton Review of Inequalities*, 2021.

Athey, Susan, Christopher Avery, and Peter Zemsky, "Mentoring and diversity," *American Economic Review*, 2000, *90* (4), 765–786.

Bertrand, Marianne, "Gender in the twenty-first century," in "AEA Papers and proceedings," Vol. 110 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2020, pp. 1–24.

_ **and Esther Duflo**, "Field experiments on discrimination," *Handbook of economic field experiments*, 2017, *1*, 309–393.

Besley, Timothy, Olle Folke, Torsten Persson, and Johanna Rickne, "Gender Quotas and the Crisis of the Mediocre Man: Theory and Evidence from Sweden," *American Economic Review*, August 2017, *107* (8), 2204–42.

Blau, Francine D, *Equal pay in the office*, Lexington Books, 1977.

Borjas, George J., "Self-Selection and the Earnings of Immigrants," *The American Economic Review*, 1987, *77* (4), 531–553.

Bursztyn, Leonardo, Alexander W Cappelen, Bertil Tungodden, Alessandra Voena, and David H Yanagizawa-Drott, "How Are Gender Norms Perceived?," *National Bureau of Economic Research*, 2023.

Card, David, Ana Rute Cardoso, and Patrick Kline, "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women*," *The Quarterly Journal of Economics*, 10 2015, *131* (2), 633–686.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde, "Global evidence on economic preferences," *The Quarterly Journal of Economics*, 2018, *133* (4), 1645–1692.

__ , __ , **Thomas J Dohmen, David Huffman, and Uwe Sunde**, "The preference survey module: A validated instrument for measuring risk, time, and social preferences," 2016.

Fernández, Raquel, Alessandra Fogli, and Claudia Olivetti, "Mothers and sons: Preference formation and female labor force dynamics," *The Quarterly Journal of Economics*, 2004, *119* (4), 1249–1299.

__ **, Asel Isakova, Francesco Luna, and Barbara Rambousek**, *Gender equality and inclusive growth*, International Monetary Fund, 2021.

Goldin, Claudia, "The U-shaped female labor force function in economic development and economic history," 1995.

__ , "A grand gender convergence: Its last chapter," *American Economic Review*, 2014, *104* (4), 1091–1119.

Hamilton, Barton H, Jack A Nickerson, and Hideo Owan, "Diversity and productivity in production teams," in "Advances in the Economic Analysis of participatory and Labor-managed Firms," Vol. 13, Emerald Group Publishing Limited, 2012, pp. 99–138.

Heckman, James, "The Common Structure of Statistical Models of Truncation, Sample Selection and Limited Dependent Variables and a Simple Estimator for Such Models," in "Annals of Economic and Social Measurement," Vol. 5, National Bureau of Economic Research, Inc, 1976, pp. 475–492.

Hjort, Jonas, "Ethnic divisions and production in firms," *The Quarterly Journal of Economics*, 2014, *129* (4), 1899–1946.

__ **, Xuan Li, and Heather Sarsons**, "Across-Country Wage Compression in Multinationals," Technical Report, National Bureau of Economic Research 2020.

Hong, Lu and Scott E Page, "Problem solving by heterogeneous agents," *Journal of economic theory*, 2001, *97* (1), 123–163.

Hsieh, Chang-Tai, Erik Hurst, Charles I Jones, and Peter J Klenow, "The allocation of talent and us economic growth," *Econometrica*, 2019, *87* (5), 1439–1474.

Hyland, Marie, Simeon Djankov, and Pinelopi Koujianou Goldberg, "Do Gendered Laws Matter?," 2021.

Jayachandran, Seema, "The roots of gender inequality in developing countries," *economics*, 2015, *7* (1), 63–88.

Lazear, Edward P, "Globalisation and the market for team-mates," *The Economic Journal*, 1999, *109* (454), 15–40.

_ , "Why are some immigrant groups more successful than others?," *Journal of Labor Economics*, 2021, *39* (1), 115–133.

Marx, Benjamin, Vincent Pons, and Tavneet Suri, "Diversity and team performance in a Kenyan organization," *Journal of Public Economics*, 2021, *197*, 104332.

Minni, Virginia, "Making the Invisible Hand Visible: Managers and the Allocation of Workers to Jobs," *Working Paper*, 2022.

Olivetti, Claudia and Barbara Petrongolo, "Unequal Pay or Unequal Employment? A Cross-Country Analysis of Gender Gaps," *Journal of Labor Economics*, 2008, *26* (4), 621–654.

_ **and** _ , "The evolution of gender gaps in industrialized countries," *Annual review of Economics*, 2016, *8*, 405–434.

Propper, Carol and John van Reenen, "Can Pay Regulation Kill? Panel Data Evidence on the Effect of Labor Markets on Hospital Performance," *Journal of Political Economy*, 2010, *118* (2), 222–273.

Roy, Andrew Donald, "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 1951, *3* (2), 135–146.

Wiswall, Matthew and Basit Zafar, "Preference for the Workplace, Investment in Human Capital, and Gender*," *The Quarterly Journal of Economics*, 08 2017, *133* (1), 457–507.

3.8 Figures

Figure 3.1: LFPR and GDP per capita

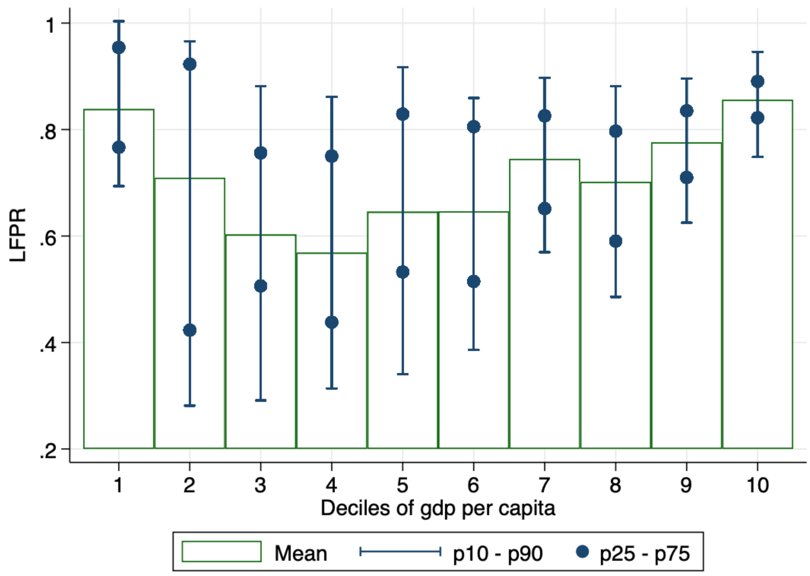

Notes. The figure plots the average, the interdecile range, and interquartile range of the LFPR across deciles of GDP per capita across countries and cohorts.

Figure 3.2: Female/Male in MNE vs LFPR

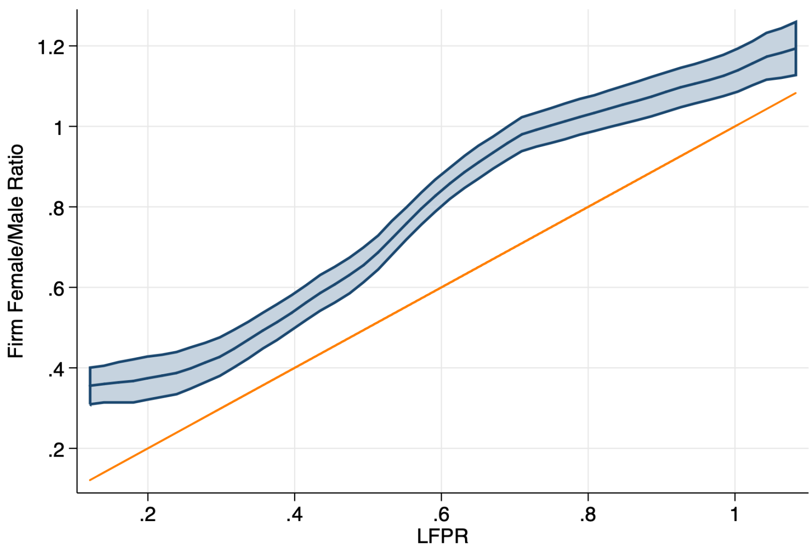

Notes. The y-axis corresponds to the female/male employment ratio in the MNE while the x-axis corresponds to the LFPR in the countries. The blue line shows the relationship between the two measures smoothed through a local linear regression. The orange line represents the 45 degree line.

Figure 3.3: Share of women and women's performance

(a) Share of women, top decile of wages

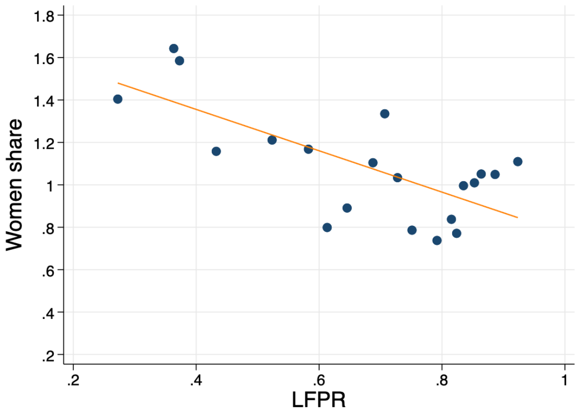

(b) Share of women, bottom decile of wages

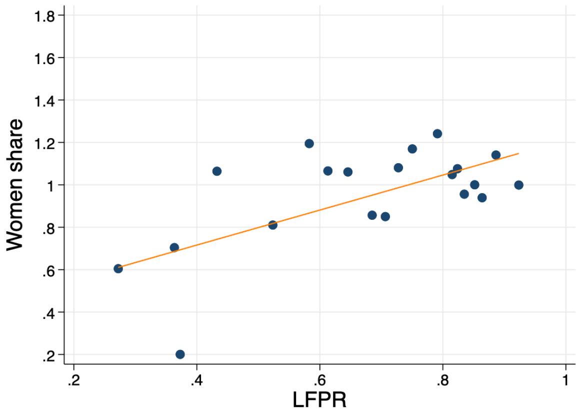

(c) Share of women, managerial positions

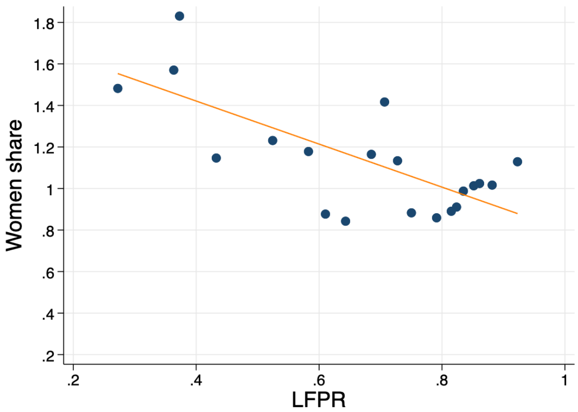

(d) Share of women, high promotion speed

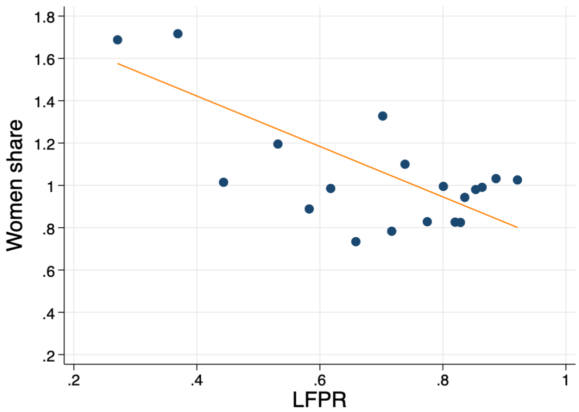

Notes. The figures are binned scatterplots and a linear fit of the share of women with different performance characteristics as a proportion of the gender share in each country-cohort cell against the LFPR. Panel (a) looks at the top decile of wages; Panel (b) at the bottom decile; Panel (c) looks at the share of women in managerial positions; and Panel (d) at promotion speed to managerial positions (we apply the method developed in Minni (2022) to identify fast managerial promotions). In the regressions, we use analytical weights by employee size of each cohort-country cell.

Figure 3.4: Women's wages and the LFPR

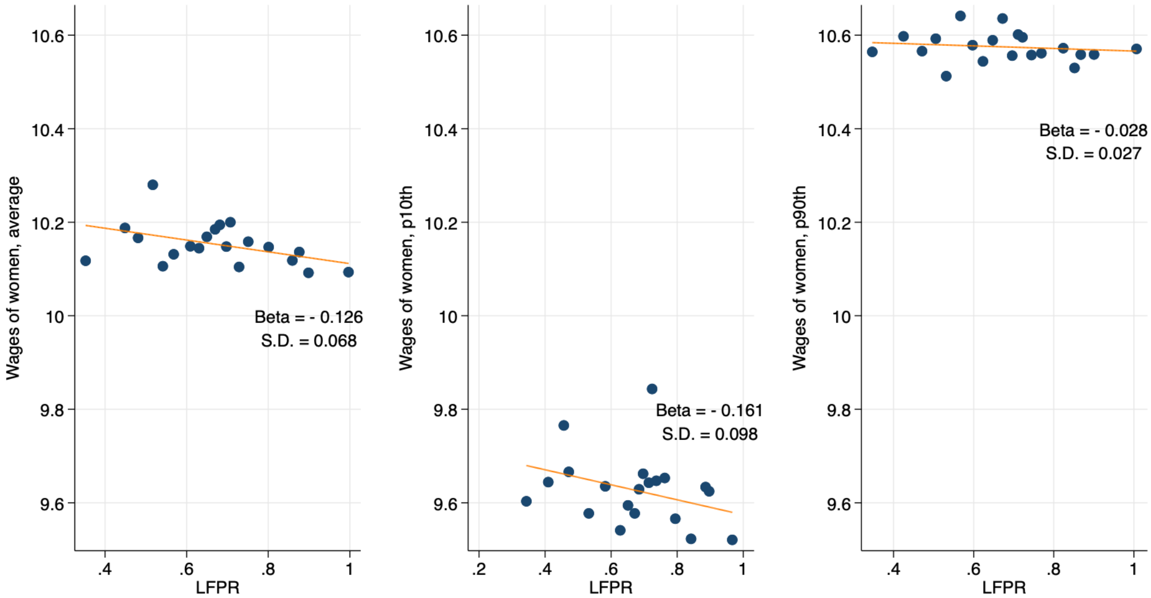

Notes. The figures are binned scatterplots and a linear fit of women's wages against the LFPR. The first figure from the left plots women's average wages, the middle one plots the 10th percentile, and the last one plots the 90th percentile. In the regressions, we control for the respective measure of men's wages, and we use analytical weights by employee size of each cohort-country cell and robust standard errors.

Figure 3.5: Calibrated productivity: average by LFPR

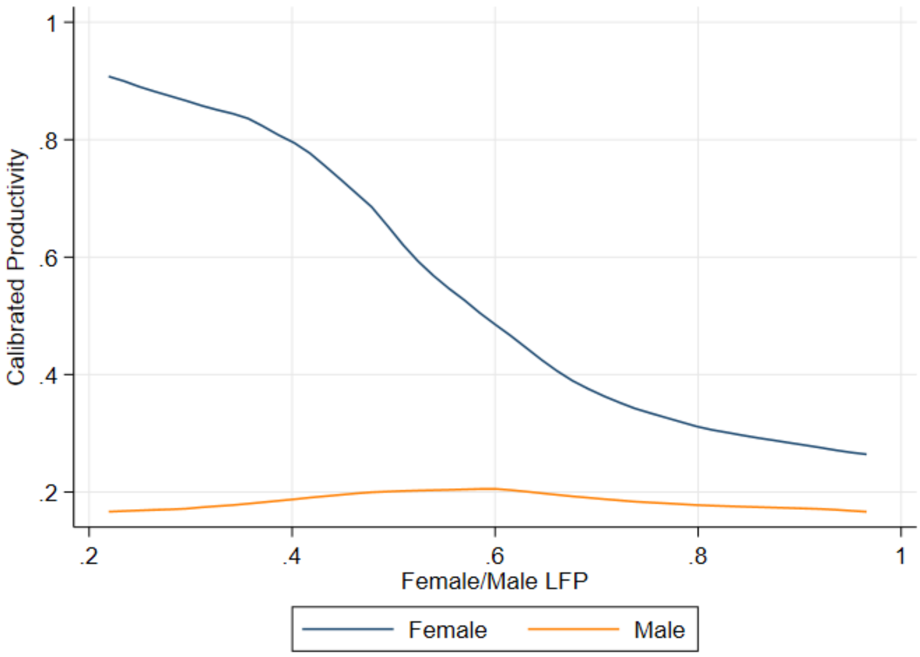

Notes. The figure plots average calibrated productivity for our sample of firm workers by Female/Male LFP, smoothed through a local linear regression.

Figure 3.6: Calibrated productivity: density by LFPR group

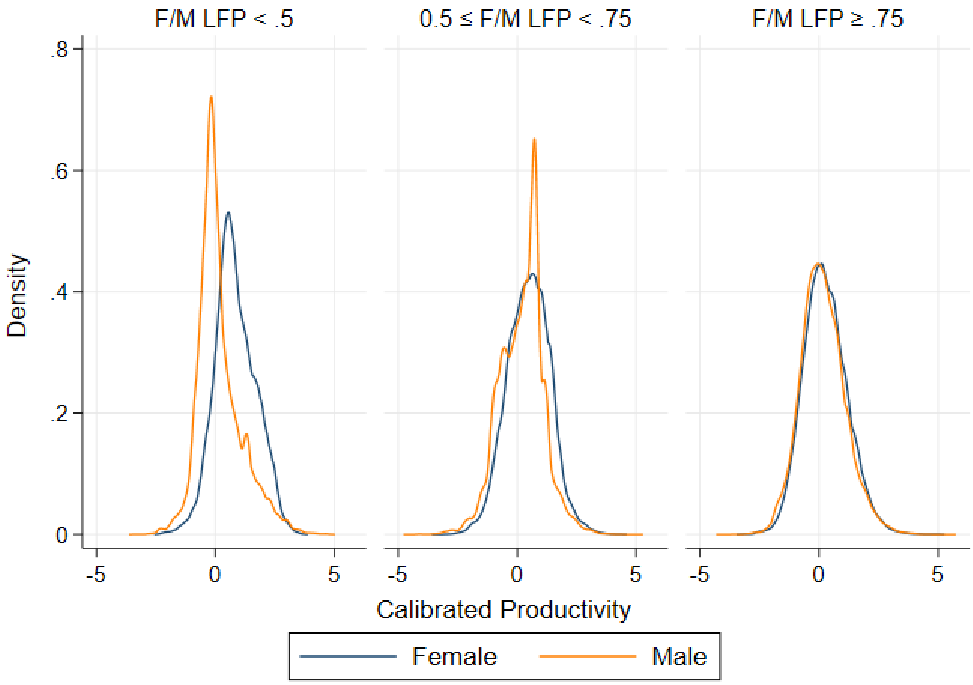

Notes. The figure plots a kernel density estimate of calibrated productivity for our sample of firm workers by three LFPR groups: $[0, .5)$, $[.5, .75)$ and $[.75, 1]$.

Figure 3.7: Counterfactual preference gap

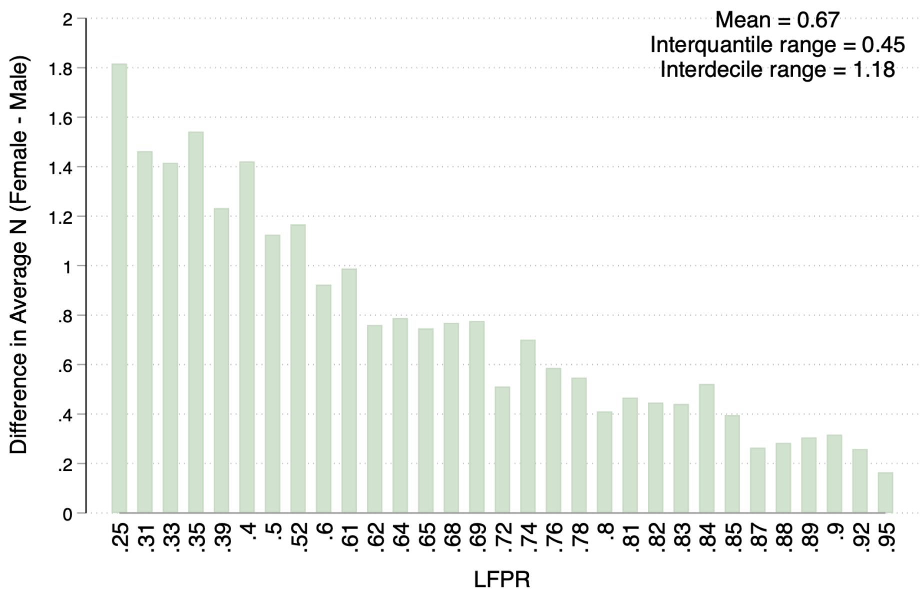

Notes. The figure plots the difference in means of N between genders that make the observed LFP gap optimal, computed as explained in the main text. The x-axis is the LFPR in each country.

Figure 3.8: Baseline vs. optimal wage policy

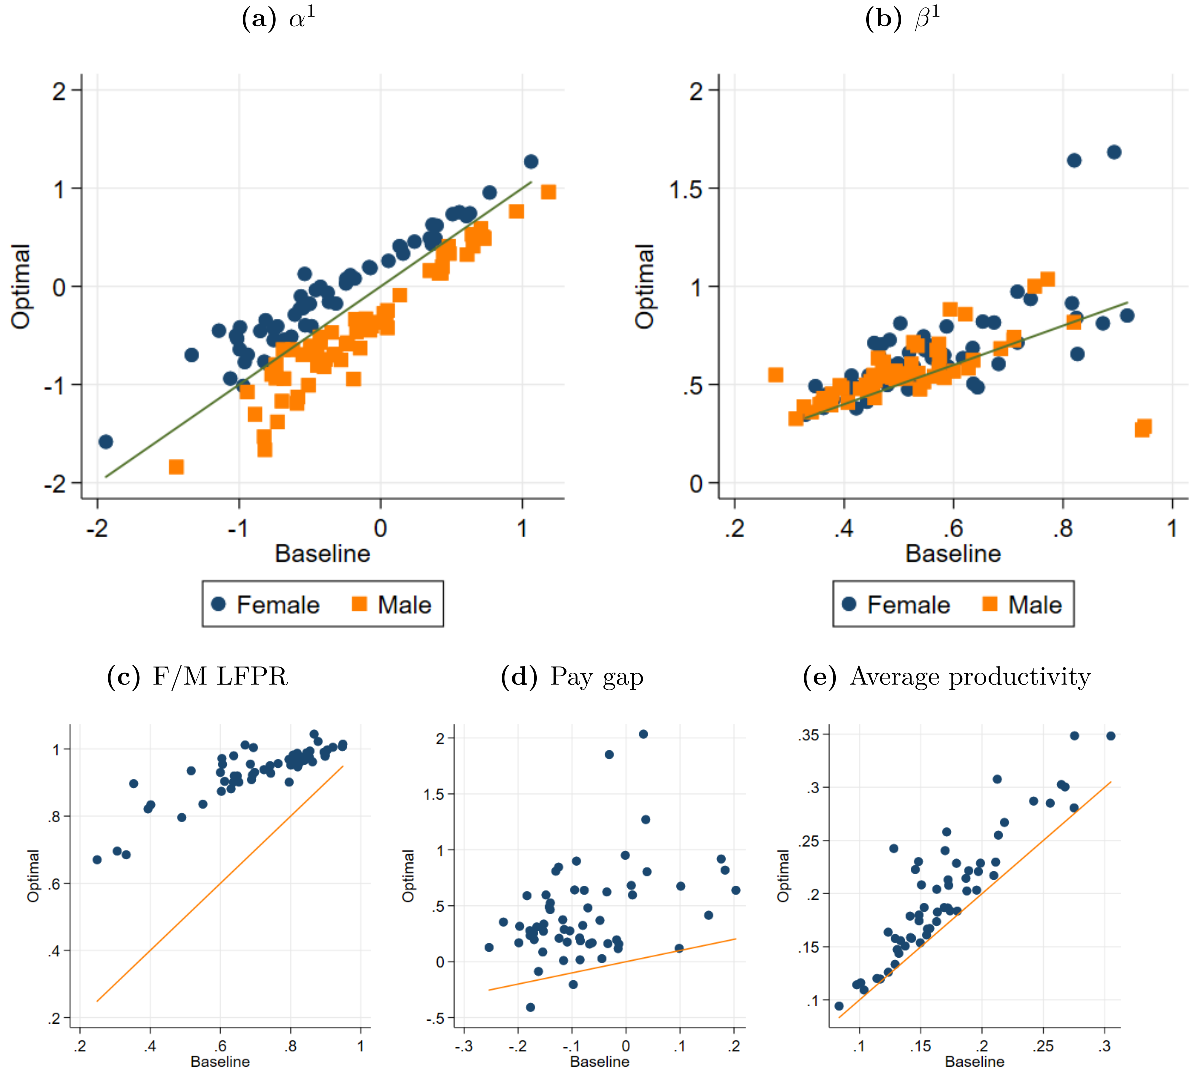

Notes. The figures compare different outcomes (female to male LFP ratio, pay gap, and average productivity) and the wage policy parameters (α^1, β^1) at baseline vs. the optimal wage policy (see main text for details). Each dot represents a country and the 45-degree line is included.

Figure 3.9: Baseline vs. equating norms (short run)

(a) F/M LFPR

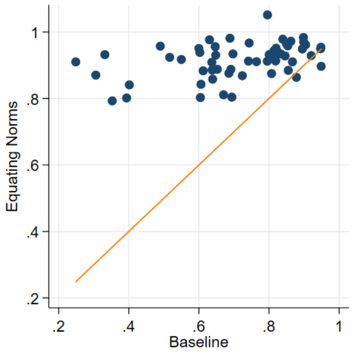

(b) Pay gap

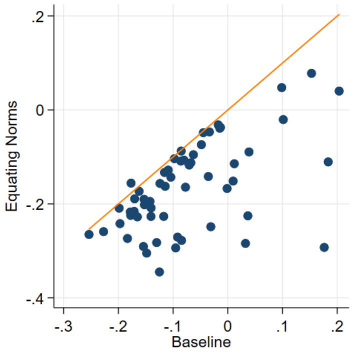

(c) Average productivity

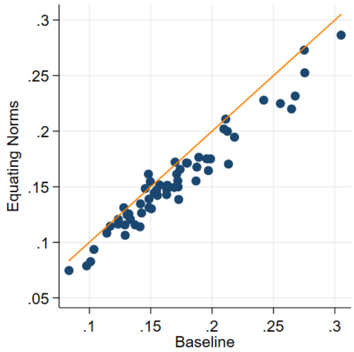

Notes. The figures compare different outcomes (female to male LFP ratio, pay gap, and average productivity) under the baseline norm parameters (α^0, ν^0) to the counterfactual where these are equalized at the male levels. The "short-run" wage policy scenario keeps the wage policy of the firm fixed at the calibrated baseline parameters. Each dot represents a country and the 45-degree line is included.

Figure 3.10: Baseline vs. equating norms (long run)

(a) α^1

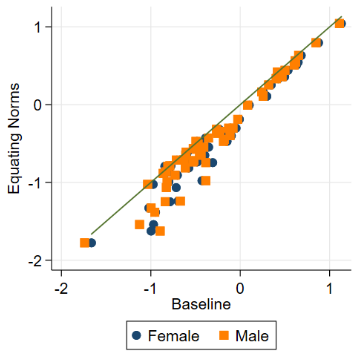

(b) β^1

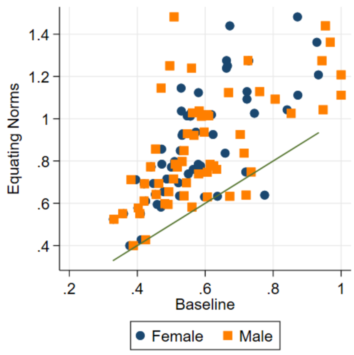

(c) F/M LFPR

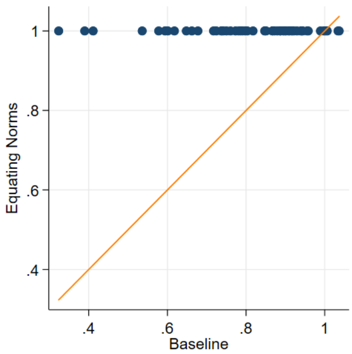

(d) Pay gap

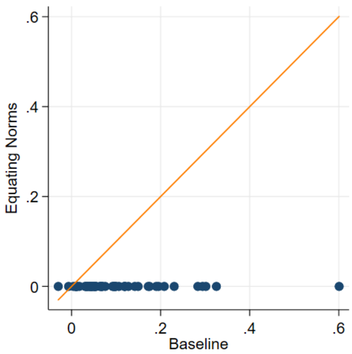

(e) Average productivity

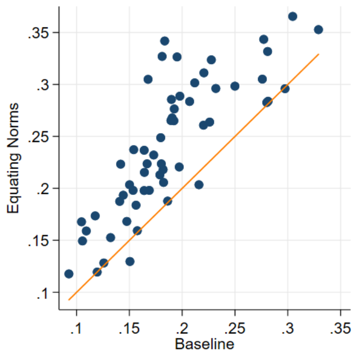

Notes. The figures compare different outcomes (female to male LFP ratio, pay gap, and average productivity) and the wage policy parameters (α^1, β^1) under the baseline norm parameters (α^0, ν^0) to the counterfactual where these are equalized at the male levels. The "long-run" wage policy scenario lets the firm optimize the wage policy to maximize productivity under certain constraints (see the main text for details). Each dot represents a country and the 45-degree line is included.

Figure 3.11: Average productivity of entrants and leavers by country (long run)

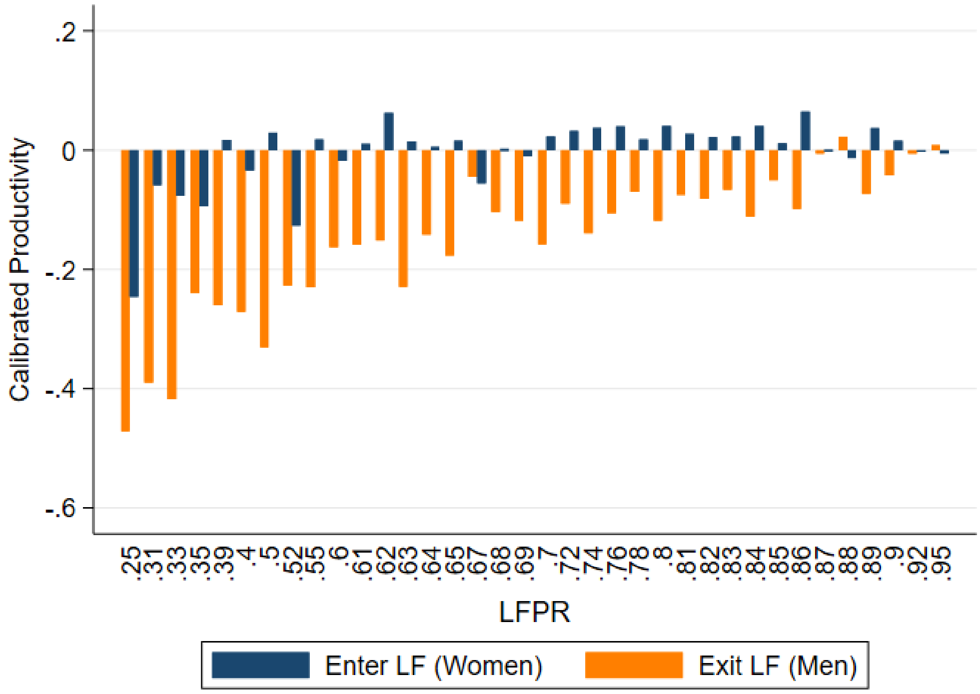

Notes. The figure plots the average productivity of LFP entrants against leavers in the long-run wage policy response. The x-axis is the LFPR in each country.

Figure 3.12: Increase in average productivity by LFPR (long run)

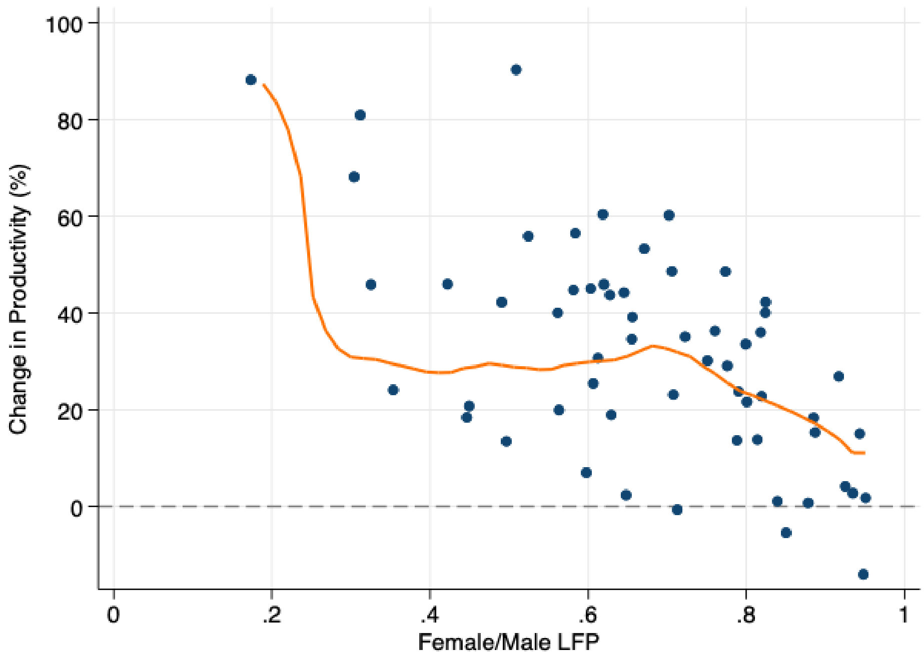

Notes. The figure plots the percentage increase in average productivity when eliminating social norms. Each dot is a country. The average across countries is 32%.

Figure 3.13: Counterfactual average productivity with strict labor regulation

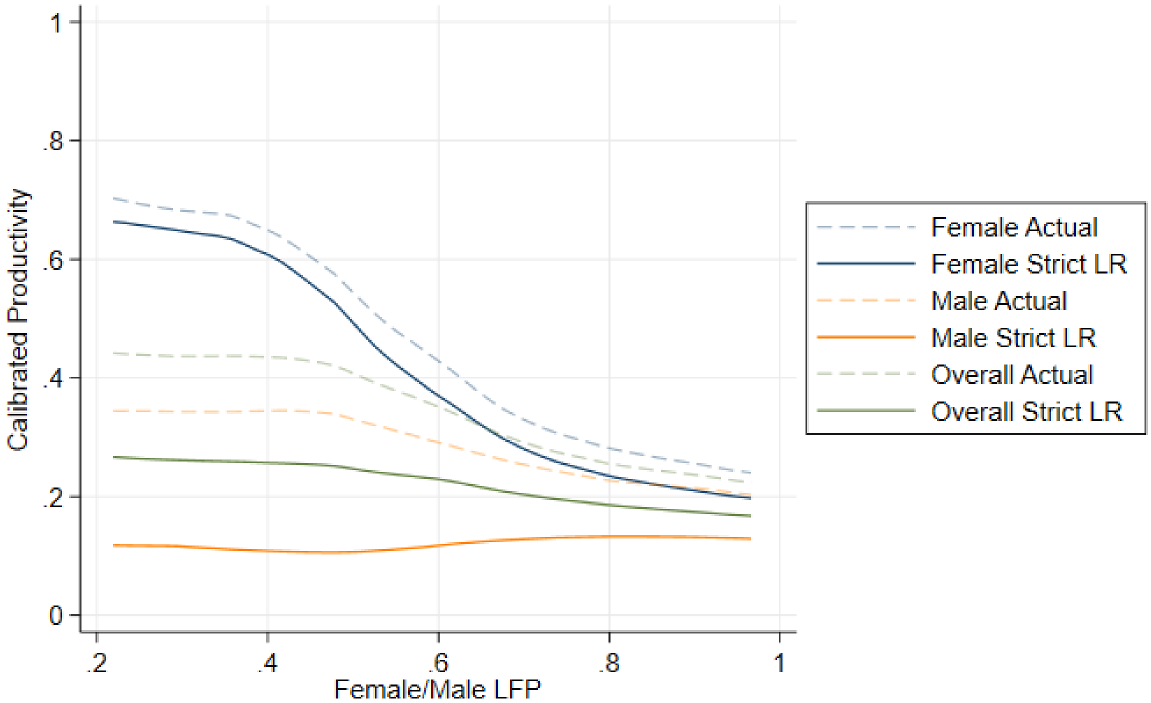

Notes. The figure plots a local polynomial regression of average productivity at baseline and under the strict labor regulation counterfactual, both estimated under the constrained optimal policy.

3.9 Tables

Table 3.1: Summary statistics

Variable	(1) Male	(2) Female
Pay + Bonus (logs)	10.467	10.426
	(0.695)	(0.648)
Age	42.859	42.404
	(10.074)	(10.159)
Tenure	12.027	11.869
	(8.605)	(8.609)
Share in Work-level 2+	0.243	0.214
	(0.286)	(0.272)
Share with fast promotions	0.186	0.180
	(0.305)	(0.301)
Share top performers	0.137	0.124
	(0.205)	(0.189)
Econ, Business, and Admin	0.506	0.526
	(0.381)	(0.371)
Share in Sales Function	0.461	0.335
	(0.337)	(0.316)
Observations	3,338	3,103

Notes. This table reports summary statistics for the relevant sample of workers used in the analysis. An observation is a gender-cohort-country-tenure cell (tenure is binned in groups of 2 years each). This is the relevant unit in the structural estimation. Tenure is measured in years. Work level denotes the hierarchical tier (from level 1 at the bottom to level 6). The share of fast promotions only considers workers that achieve at least work-level 2 or higher. The sales function is the most common function (39%). Top performers based on firm performance appraisals system (top 10%).

Table 3.2: Gender pay gap and LFPR

	Full Sample					New Hires	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			Pay + Bonus (logs)			**Pay Growth**	**Major promotion**
Female	0.377	0.256	0.028	0.253	0.197	0.159	0.107
	(0.142)	(0.096)	(0.173)	(0.094)	(0.077)	(0.067)	(0.047)
LFPR	1.640	1.641	0.596	1.662	0.138	0.514	-0.028
	(0.282)	(0.212)	(0.153)	(0.204)	(0.235)	(0.151)	(0.080)
Female × LFPR	-0.564	-0.471	-0.274	-0.451	-0.376	-0.142	-0.112
	(0.194)	(0.135)	(0.116)	(0.132)	(0.103)	(0.087)	(0.062)
GDP per capita (logs)			0.275				
			(0.021)				
Female × GDP per capita (logs)			0.008				
			(0.016)				
Controls	No	Yes	Yes	Yes	Yes	No	No
Cohort FE	No	No	No	Yes	No	No	No
Country FE	No	No	No	No	Yes	No	No
N	303756	303756	302567	303756	303756	8274	8274
R-squared	0.116	0.285	0.435	0.307	0.540	0.103	0.056

Notes. An observation is a worker-year. Controls include: tenure, tenure squared, year FE and function FE. The last two columns report estimates when restricting the sample to new hires at the entry level observed for at least four years. *Pay growth* is computed as the difference in log pay between the last year a worker is observed and the first year a worker is observed. *Probability of promotion* equals 1 if the worker was promoted to work-level 2 during the sample period. Controlling for starting salary in columns 6 and 7. Standard errors clustered at the country-cohort level.

Table 3.3: Summary of model parameters and empirical targets

Param.	Interpretation	Empirical Target
α^0_{gtac}	Unconditional average value of staying at home	LFP
α^1_{gtac}	Unconditional average log-wage	Average observed log-wage (controlling for selection)
β^1_{gtac}	Returns to productivity in the firm	Variance of the observed log-wage (controlling for selection)
ν^0_{gtac}	Dispersion of the idiosyncratic taste for staying at home	Not identified separately (Normalize $(\beta^1_{gtac})^2 + (\nu^0_{gtac})^2 = 1$)

Table 3.4: $\overline{A}_F - \overline{A}_M$ and log(operating revenue), ORBIS

	Log(OpRev)		
	(1)	(2)	(3)
Productivity gap	-9.151	-12.975	
	(1.281)	(1.879)	
Same SIC 3=1 × Productivity gap			-0.165
			(0.039)
Log(employment)	0.607	0.616	0.610
	(0.212)	(0.215)	(0.001)
Same SIC 3=1 × Log(employment)			0.148
			(0.002)
Log(capital)	0.425	0.420	0.372
	(0.061)	(0.059)	(0.001)
Same SIC 3=1 × Log(capital)			0.016
			(0.001)
Log(GDP)	0.532	0.517	
	(0.156)	(0.136)	
LFPR		-6.530	
		(3.357)	
Same SIC 3=1			-0.363
			(0.014)
Country FE	No	No	Yes
R-squared	0.665	0.672	0.714
N	2239881	2239881	2239881

Notes. An observation is a firm in the Orbis database. Cross-section based on latest year up to 2019, sample restricted to firms whose latest year is after 2011. Standard errors clustered at the country level in cols. 1 and 2 and robust in col. 3.

Table 3.5: $\overline{A}_F - \overline{A}_M$ and productivity dispersion, ORBIS

	Log(OpRev/emp.), Mean			**Log(OpRev/emp.), CV**		
	(1)	(2)	(3)	(4)	(5)	(6)
Productivity gap	-17.549	-11.390	-11.310	0.577	0.556	0.529
	(3.757)	(2.422)	(2.492)	(0.168)	(0.144)	(0.145)
Log(GDP)	0.868	0.483	0.483	0.007	0.011	0.011
	(0.248)	(0.173)	(0.173)	(0.008)	(0.009)	(0.009)
LFPR	-7.750	-6.892	-6.907	0.365	0.388	0.386
	(5.438)	(3.493)	(3.501)	(0.278)	(0.242)	(0.242)
Log(capital), Mean		0.458	0.449			
		(0.065)	(0.065)			
Same SIC 3=1 $\times$ Productivity gap			-0.379			0.143
			(0.613)			(0.039)
Same SIC 3=1 $\times$ Log(capital), Mean			0.062			
			(0.019)			
Log(capital), CV					0.148	0.143
					(0.085)	(0.086)
Same SIC 3=1 $\times$ Log(capital), CV						0.037
						(0.032)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.751	0.890	0.891	0.485	0.523	0.529
N	2418	2418	2418	2418	2418	2418
Outcome mean	10.822	10.822	10.822	0.145	0.145	0.145
Outcome sd	2.006	2.006	2.006	0.059	0.059	0.059

Notes. An observation is an industry (US SIC 3) -country cell in the Orbis database. Analytics weights used. Measures based on cross-section of firms based on latest year up to 2019, sample restricted to firms whose latest year is after 2011. Standard errors clustered at the country level.

3.10 Appendix

3.10.1 Descriptive analysis

Figure 3.14: The countries where the MNE operates and female to male LFP

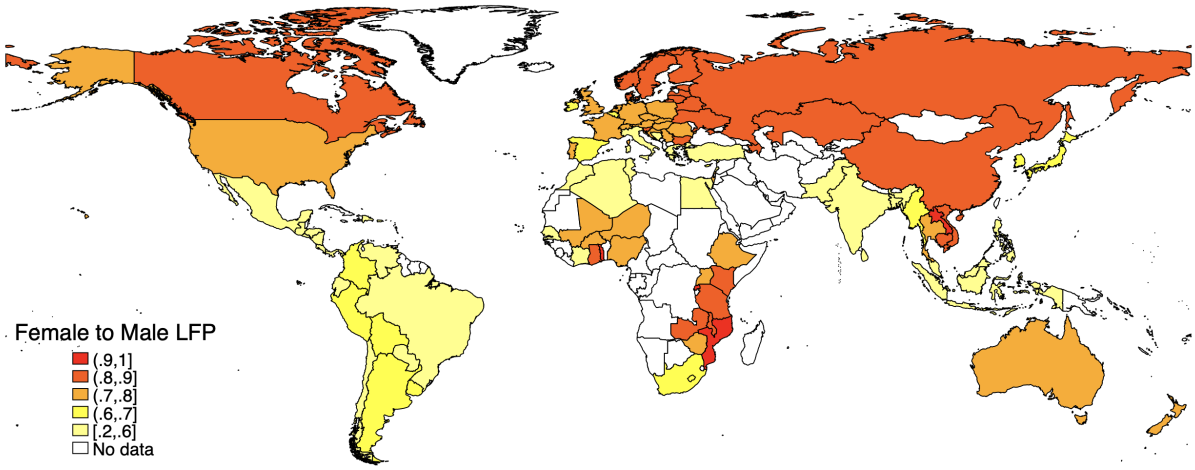

Figure 3.15: Average wages in the firm and in the country overall: a) ILO, white collar occupations only; and b) ORBIS, manufacturing sector only

(a) Wage ratios against LFPR

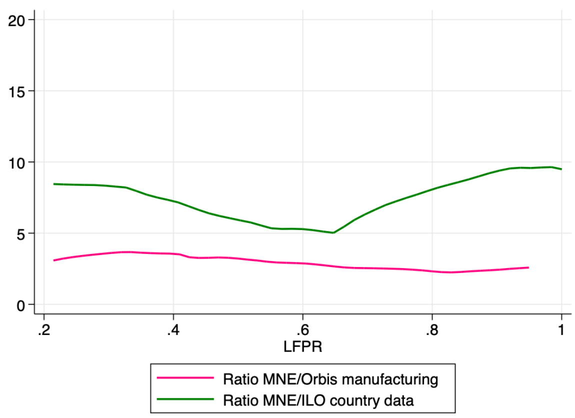

(b) Wage ratios against GDP per capita in logs

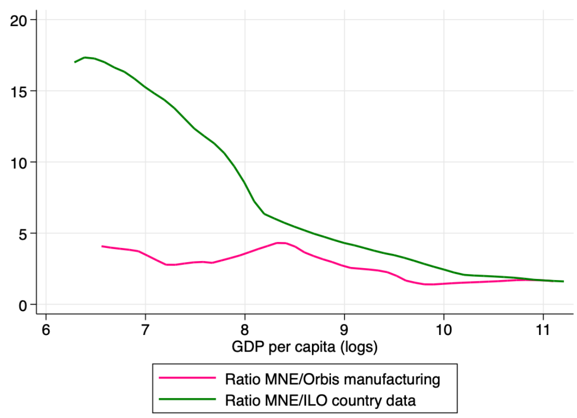

Notes. This figure plots the ratio of the average wage in the MNE and in the country overall: a) from the ORBIS database, considering the manufacturing sector only, and b) from the International Labor Organization, considering white-collar occupations only. Wages are measured in 2017 PPP \$. The x-axis is the LFPR (panel a) and the GDP per capita in logs (panel b) in each country.

Table 3.6: Gender pay gap and LFPR — LFP for population with advanced education

	Pay + Bonus (logs)				
	(1)	(2)	(3)	(4)	(5)
Female	0.427	0.269	0.161	0.303	0.381
	(0.173)	(0.082)	(0.150)	(0.079)	(0.073)
LFPR, advanced education	1.776	1.692	-0.138	1.684	0.073
	(0.188)	(0.113)	(0.222)	(0.116)	(0.273)
Female × LFPR, advanced education	-0.496	-0.375	-0.484	-0.401	-0.522
	(0.198)	(0.100)	(0.140)	(0.096)	(0.088)
GDP per capita (logs)			0.358		
			(0.033)		
Female × GDP per capita (logs)			0.020		
			(0.016)		
Controls	No	Yes	Yes	Yes	Yes
Cohort FE	No	No	No	Yes	Yes
Country FE	No	No	No	No	No
N	251336	251336	250152	251336	251336
R-squared	0.155	0.315	0.469	0.332	0.570

Notes. An observation is a worker-year. The LFPR is computed using the LFP for individuals with advanced education (short-cycle tertiary education or college degree and/or above). Controls include: tenure, tenure squared, year FE and function FE. Standard errors clustered at the country-cohort level.

Table 3.7: Gender pay gap and LFPR — by region and income group

	All	Region FE	Lower inc.	Higher inc.
	(1)	(2)	(3)	(4)
	Pay + Bonus (logs)			
Female	0.256	0.208	0.265	-0.032
	(0.096)	(0.070)	(0.050)	(0.111)
LFPR	1.641	0.698	0.547	1.384
	(0.212)	(0.186)	(0.134)	(0.338)
Female × LFPR	-0.471	-0.424	-0.408	-0.129
	(0.135)	(0.100)	(0.068)	(0.156)
Controls	Yes	Yes	Yes	Yes
N	303756	303756	71658	232098
R-squared	0.285	0.387	0.218	0.239

Notes. An observation is a worker-year. Controls include: tenure, tenure squared, year FE and function FE. Standard errors clustered at the country-cohort level. Income group and geographical region are obtained from the World Bank.

Table 3.8: Gender pay gap and LFPR — PPP conversion

	Pay + Bonus (logs), PPP 2017 USD				
	(1)	(2)	(3)	(4)	(5)
Female	0.318	0.195	0.312	0.197	0.198
	(0.127)	(0.080)	(0.141)	(0.080)	(0.076)
LFPR	0.266	0.249	0.085	0.259	0.129
	(0.200)	(0.136)	(0.155)	(0.130)	(0.239)
Female × LFPR	-0.463	-0.372	-0.275	-0.364	-0.381
	(0.165)	(0.104)	(0.083)	(0.105)	(0.102)
GDP per capita (logs)			0.044		
			(0.018)		
Female × GDP per capita (logs)			-0.021		
			(0.013)		
Controls	No	Yes	Yes	Yes	Yes
Cohort FE	No	No	No	Yes	Yes
Country FE	No	No	No	No	No
N	302789	302789	301600	302789	302789
R-squared	0.014	0.164	0.169	0.173	0.339

Notes. An observation is a worker-year. Controls include: tenure, tenure squared, year FE and function FE. Standard errors clustered at the country-cohort level. Wages are measured in PPP 2017 USD. Purchasing power parity (PPP) exchange rates are taken from the ICP (World Bank).

Table 3.9: Gender pay gap and LFPR — fixed pay only

	Pay (logs)				
	(1)	(2)	(3)	(4)	(5)
Female	0.358	0.232	-0.023	0.230	0.167
	(0.135)	(0.092)	(0.162)	(0.091)	(0.074)
LFPR	1.582	1.578	0.501	1.599	0.144
	(0.282)	(0.212)	(0.145)	(0.204)	(0.234)
Female × LFPR	-0.533	-0.436	-0.241	-0.415	-0.328
	(0.185)	(0.131)	(0.114)	(0.129)	(0.100)
GDP per capita (logs)			0.283		
			(0.020)		
Female × GDP per capita (logs)			0.011		
			(0.015)		
Controls	No	Yes	Yes	Yes	Yes
Cohort FE	No	No	No	Yes	Yes
Country FE	No	No	No	No	No
N	303756	303756	302567	303756	303756
R-squared	0.110	0.285	0.448	0.309	0.552

Notes. An observation is a worker-year. Controls include: tenure, tenure squared, year FE and function FE. Standard errors clustered at the country-cohort level.

3.10.2 Structural model

Figure 3.16: Gender differences in preferences

(a) Altruism

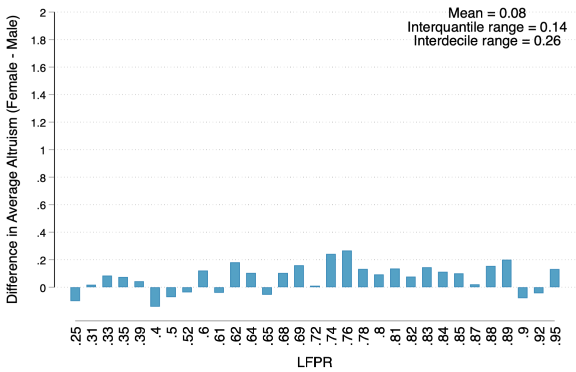

(b) Risk-taking

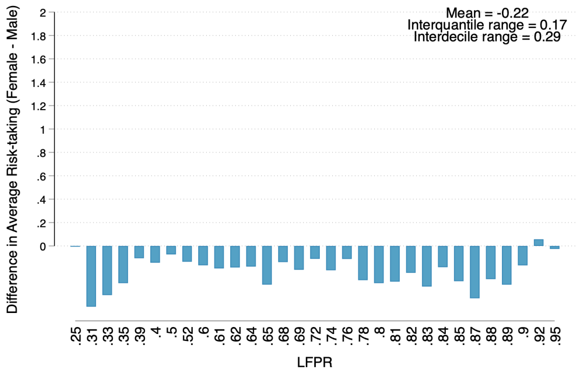

(c) Trust

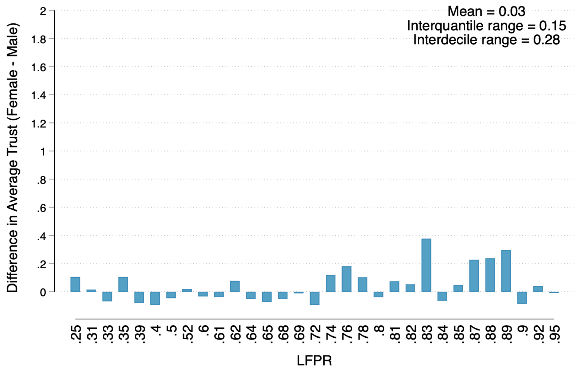

(d) Patience

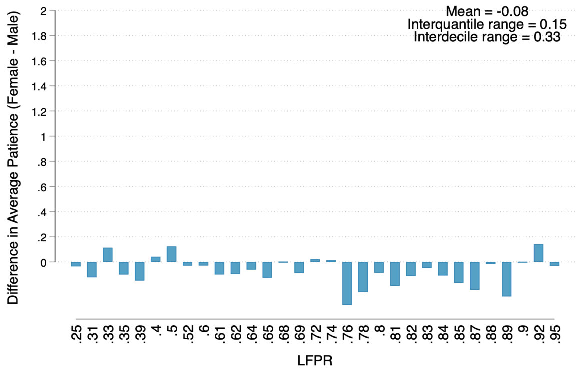

(e) Positive reciprocity

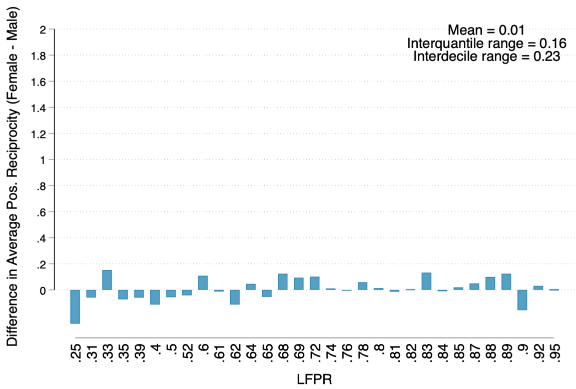

(f) Negative reciprocity

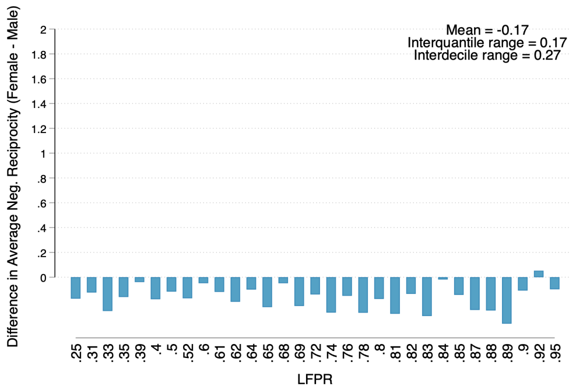

Notes. The figure plots the difference in means of economic preferences between genders. Data taken from the Global Preferences Survey (Falk et al. (2016), Falk et al. (2018)). The x-axis is the LFPR in each country.

Figure 3.17: Validation of our calibrated productivity against other performance measures

(a) Performance appraisals

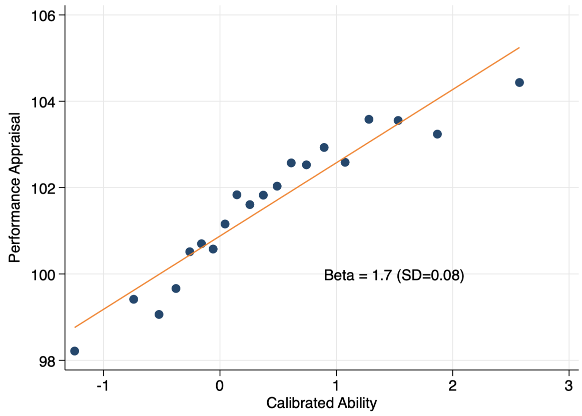

(b) Pay growth in the first year (new hires)

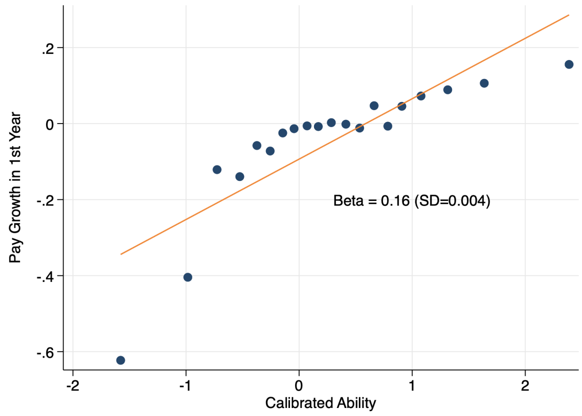

(c) Productivity (sales)

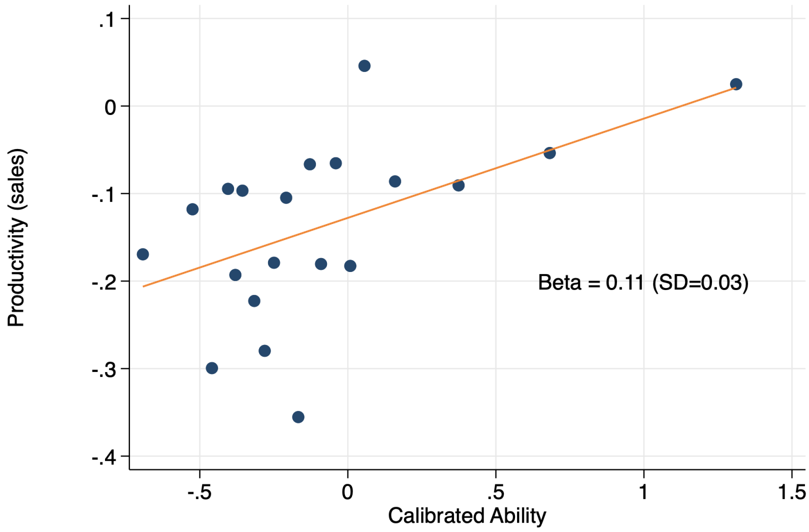

Notes. The figures are binned scatterplots and a linear fit of other performance measures (performance appraisals, pay growth for new hires, and objective productivity) against our calibrated productivity. The objective productivity measure is available only for the sales function in 10 countries and is based on reaching set targets.

Figure 3.18: Validation of our calibrated parameters against labor regulations

(a) Performance rewards (β^1) against the Restrictive Labor Regulations Index

(b) Value of staying at home ($\alpha_F^0 - \alpha_M^0$) against the Women, Business and the Law Index

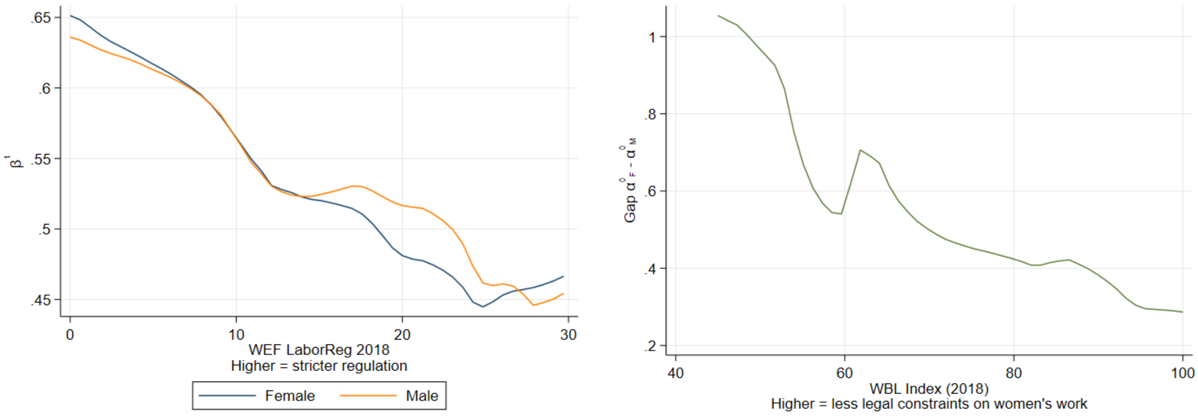

Notes. Both panels show local polynomials of our calibrated parameters against two indices related to labor regulations. Panel (a) plots the calibrated β^1_{gtac}, for men and women separately, against the WEF Restrictive Labor Regulations Index. Panel (b) plots the gap in calibrated $\alpha_F^0 - \alpha_M^0$ against the WB Women, Business, and the Law index. Details about these indices are in the main text.

Figure 3.19: Validation of our calibrated parameters against values

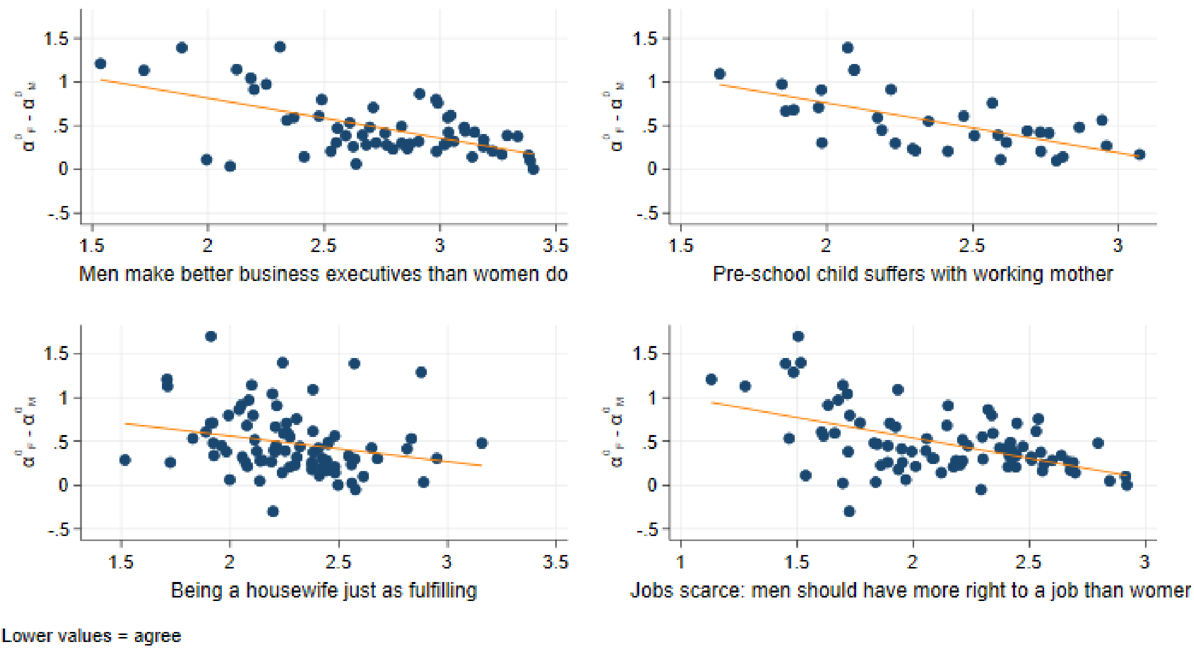

Notes. The figure shows scatterplots and fitted linear regressions of the gap in calibrated $\alpha^0_{Ftac} - \alpha^0_{Mtac}$ against four questions in the World Value Survey: (1) "Men make better business executives than women do," (2) "Pre-school child suffers with working mother," (3) "Being a housewife is just as fulfilling as working," (4) "When jobs are scarce, men should have more right to a job than women." For all questions, lower values of the index denote more agreement with the statement. Each dot is a country-cohort pair.

Figure 3.20: The gender education gap versus the LFP gap

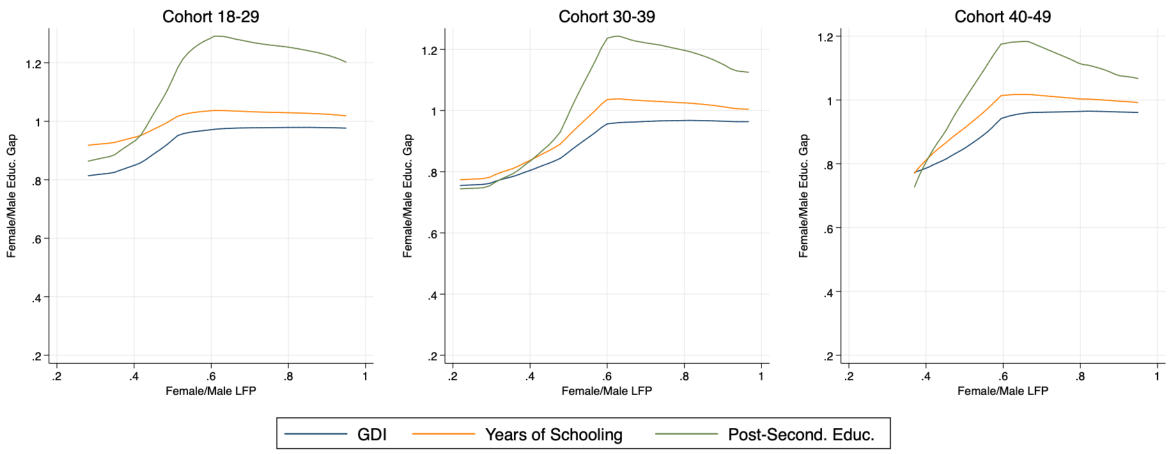

Notes. The figure plots the gender gap in education against the gender gap in LFP. We use a number of education measures: the gender development index (GDI, the ratio of female/male Human Development Index); years of schooling (female to male ratio), and the percentage in post-secondary education (female to male ratio). The GDI data is from the UNDP and the educational attainment data is from the World Bank.

3.10.3 Counterfactuals

Optimization problems for the firm's wage policy under the counterfactuals

The optimal policy we consider in subsection 3.5.1 solves, for each country-cohort-tenure cell, the following program:

$$\max_{(\alpha^1_{gtac},\beta^1_{gtac})_{g\in\{F,M\}}} \left[1-\Phi\left(\widetilde{\xi}_{Ftac}\right)\right]\frac{\beta^1_{Ftac}}{\sigma_{Ftac}}\lambda\left(\widetilde{\xi}_{Ftac}\right)+\left[1-\Phi\left(\widetilde{\xi}_{Mtac}\right)\right]\frac{\beta^1_{Mtac}}{\sigma_{Mtac}}\lambda\left(\widetilde{\xi}_{Mtac}\right)$$

$$\text{subj. to:} \quad 1-\Phi\left(\widetilde{\xi}_{Ftac}\right)+1-\Phi\left(\widetilde{\xi}_{Mtac}\right)=FLFP_{tac}+MLFP_{tac} \tag{i}$$

$$\left[1-\Phi\left(\widetilde{\xi}_{Ftac}\right)\right]\left[\alpha^1_{Ftac}+\frac{(\beta^1_{Ftac})^2}{\sigma_{Ftac}}\lambda\left(\widetilde{\xi}_{Ftac}\right)\right]+ + \left[1-\Phi\left(\widetilde{\xi}_{Mtac}\right)\right]\left[\alpha^1_{Mtac}+\frac{(\beta^1_{Ftac})^2}{\sigma_{Mtac}}\lambda\left(\widetilde{\xi}_{Mtac}\right)\right]=\overline{y}^1_{Ftac}+\overline{y}^1_{Mtac} \tag{ii}$$

where $\widetilde{\xi}_{gtac}=(\alpha^0_{gtac}-\alpha^1_{gtac})/\sigma_{gtac}$, $\sigma_{gtac}=\sqrt{(\beta^1_{gtac})^2+(\nu^0_{gtac})^2}$. The objective is average productivity in the firm ($LFP\cdot\mathbb{E}[A\,|\,\text{empl'd}]$). Constraint (i) states that total employment (or LFP) should stay constant. Constraint (ii) states that the total wage bill should be unchanged.

The long-run constrained optimal policy considered in the counterfactuals in subsection 3.5.2 solves the program above, with the additional constraint:

$$\left|\left[\alpha^1_{Ftac}+\frac{(\beta^1_{Ftac})^2}{\sigma_{Ftac}}\lambda\left(\widetilde{\xi}_{Ftac}\right)\right]-\left[\alpha^1_{Mtac}+\frac{(\beta^1_{Mtac})^2}{\sigma_{Mtac}}\lambda\left(\widetilde{\xi}_{Mtac}\right)\right]\right|\leq B \tag{iii}$$

Constraint (iii) places a limit B on how unequal the wage of the average men and women in the firm can be. The bound B is chosen to maximize the goodness of fit to the data, measured by the distance $\|\theta^*_B - \widehat{\theta}\|$ (where $\widehat{\theta}$ is the vector of calibrated $\alpha^1_{gtac}, \beta^1_{gtac}$ and θ^*_B is the solution of the constrained optimal policy problem with bound B).

Effects on workers' welfare

The effect of eliminating gender norms on wages (and hence on workers' utility) depends on the assumptions about the wage policy of the firm. If the wage parameters stay fixed at baseline levels, the only welfare effects are on the women who enter the labor force. By a revealed preference argument, those women gain because their wage in the firm is higher than their underlying value of staying at home (which is assumed to be the same as men's); the only reason why they were not in the labor force before at the current wages is the gender norm tax. There is no welfare effect on male workers nor on the women who are already in the labor force since their wages are unchanged. Hence, the net welfare effect on workers is positive in the short run.

When the firm adjusts, in contrast, because we are imposing the constraint that the wage bill stays fixed, there are winners and losers of the policy, and men can be affected by the elimination of the gender norm tax too. In particular, by the same revealed preference argument as above, those who enter the labor force will gain, and those who exit the labor force will lose. The sign of the welfare change for those who are always in the labor force depends on how the firm optimally reacts to the elimination of gender norms: in particular, if β_1 increases, as we find in Figure 3.10 panel (b), higher-productivity workers will benefit from a wage increase and lower-productivity workers will suffer from a wage decrease.

To illustrate, we focus on the two countries on opposite ends of the LFPR distribution: India and Sweden. The solid lines in Figure 3.21 represent the participation frontier, $\beta^1 A_i - \nu^0 N_i \geq \alpha^0 - \alpha^1$, for India and Sweden in the long-run horizon (averaging across cohort and tenure groups, weighted by size). The dashed lines in the same figure show the change after the gender norm tax is eliminated. LFP and the participation frontier are very equal in Sweden, to begin with, hence the effect of equating α^0, ν^0 for men and women will be very small, and effectively there will be no welfare change. In contrast, in India, the participation frontiers are very different, to begin with, and equating the value of staying at home leaves room to attract high-productivity females into the labor force. Under the long-run wage policy we consider, once the gender norm tax is eliminated, the firm can maximize productivity by equating the wage parameters for men and women (which results, on average, in a wage increase for women and a wage decrease for men). The magnitude of this change depends on how constrained the firm was originally, as shown in Figure 3.22.

The results for all countries are in Figure 3.23, which plots the change in the value of chosen option y^* (which can be interpreted as a percentage change, because it is in units of log-wages). For men, $y^* = \max\{y^0, y^1\}$. For women, $y^* = y^1_F \mathbf{1}\{y^1_F \geq y^0_F\} + y^0_M \mathbf{1}\{y^1_F < y^0_F\}$ (i.e. they make a decision based on y^0_F but the true value of staying at home is y^0_M, since we interpret the difference $y^0_F - y^0_M$ as the gender norm tax).

Whether it is men or women who lose depends strongly on how much flexibility the firm had to react to differences in gender norms in the first place, and it is best understood by looking at the parameters under the counterfactual policies in Figure 3.10. For women, the increase in LFP and in variable pay is big enough to offset the decrease in fixed pay, so they are net winners. Men's LFP decreases and fixed pay falls enough that their overall welfare change is negative. Because we impose the constraint that the wage bill is kept constant, welfare changes are essentially a redistribution between female and male workers. It is important to note that this result is driven by the assumption that the firm cannot adjust employment levels. This is convenient for quantifying the pure effect of misallocation but it does not capture all the possible benefits of expanding women's access to the labor force. In a more general model where the firm can adjust employment in response to the change in gender norms, the effect is unlikely to be zero-sum.

Figure 3.21: Examples: India and Sweden before and after equating norms

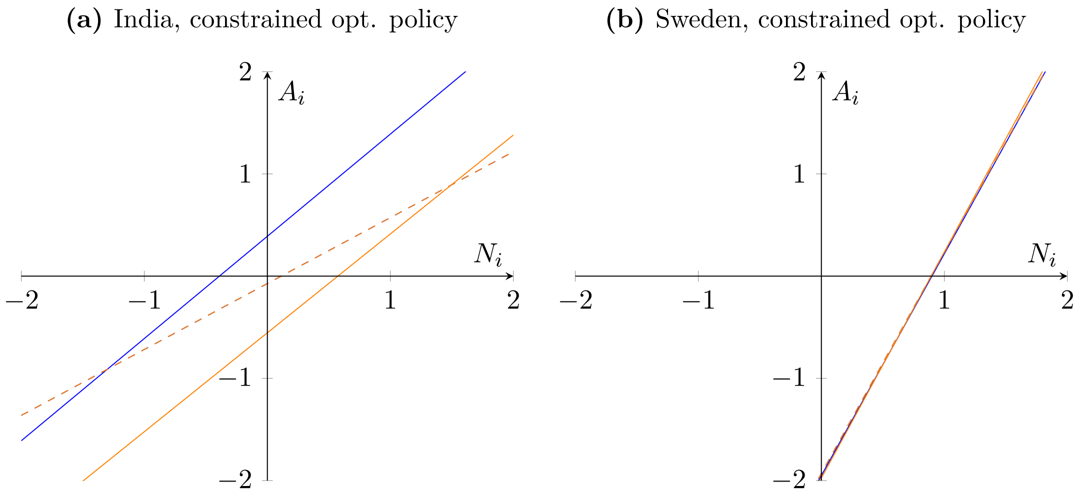

Notes. The figures show the participation frontier, $\beta^1 A_i - \nu^0 N_i \geq \alpha^0 - \alpha^1$, for India and Sweden in the long-run horizon (averaging across cohort and tenure groups, weighted by size). Blue represents females, orange represents males. The solid line is the participation frontier with baseline α^0, ν^0; the dashed line is the counterfactual participation frontier after equating gender norms (and letting the wage policy respond). Those with (A_i, N_i) above and to the left of the participation frontier are in the labor force.

Figure 3.22: Example: India, winners, and losers under different wage policies

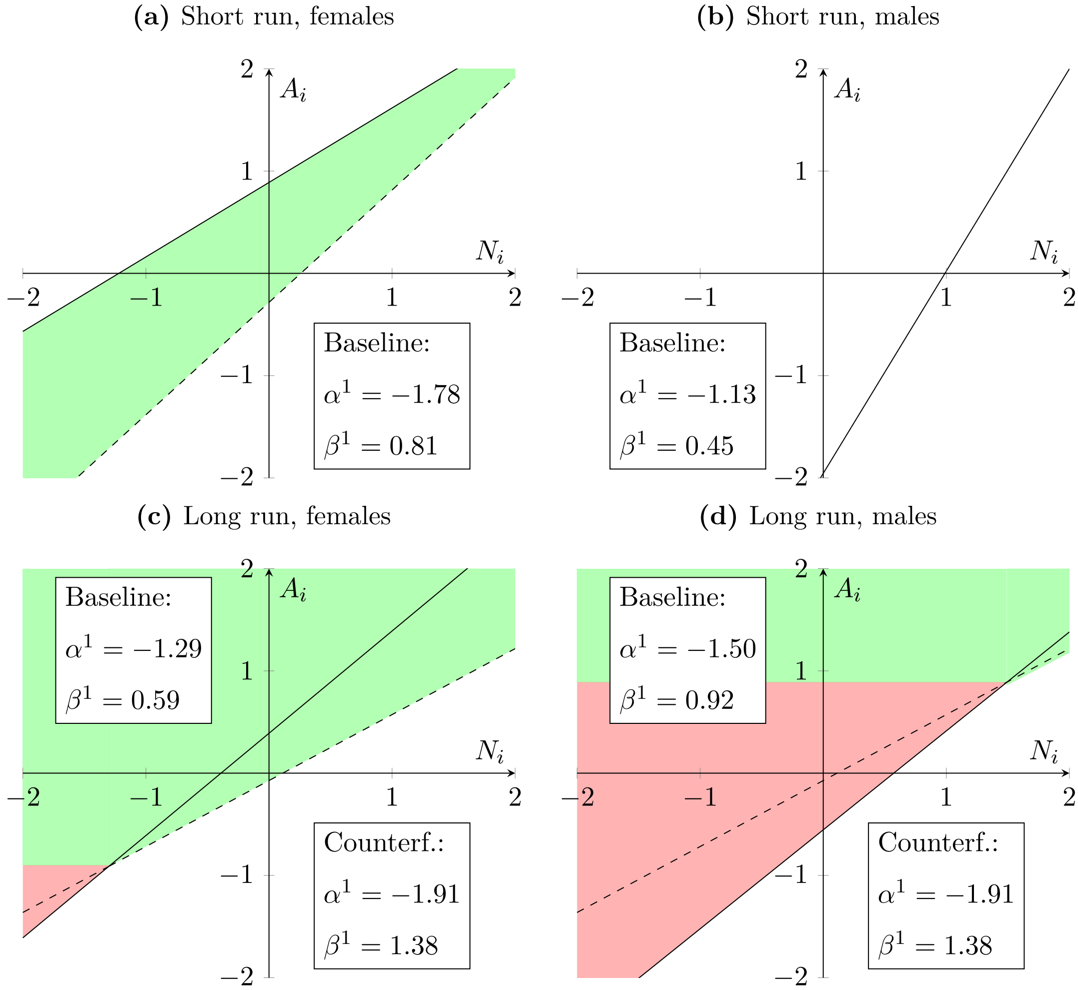

Notes. The figures show the participation frontier, $\beta^1 A_i - \nu^0 N_i \geq \alpha^0 - \alpha^1$, as well as the welfare effects for India under the two wage policy horizons we consider. The solid line is the participation frontier with baseline α^0, ν^0; the dashed line is the counterfactual participation frontier after equating gender norms (and letting the wage policy respond under the optimal policies). The area shaded green represents the population with (A_i, N_i) that benefits from equating gender norms, and the area shaded red is the population that is harmed.

Figure 3.23: Welfare effects of eliminating the gender norms tax

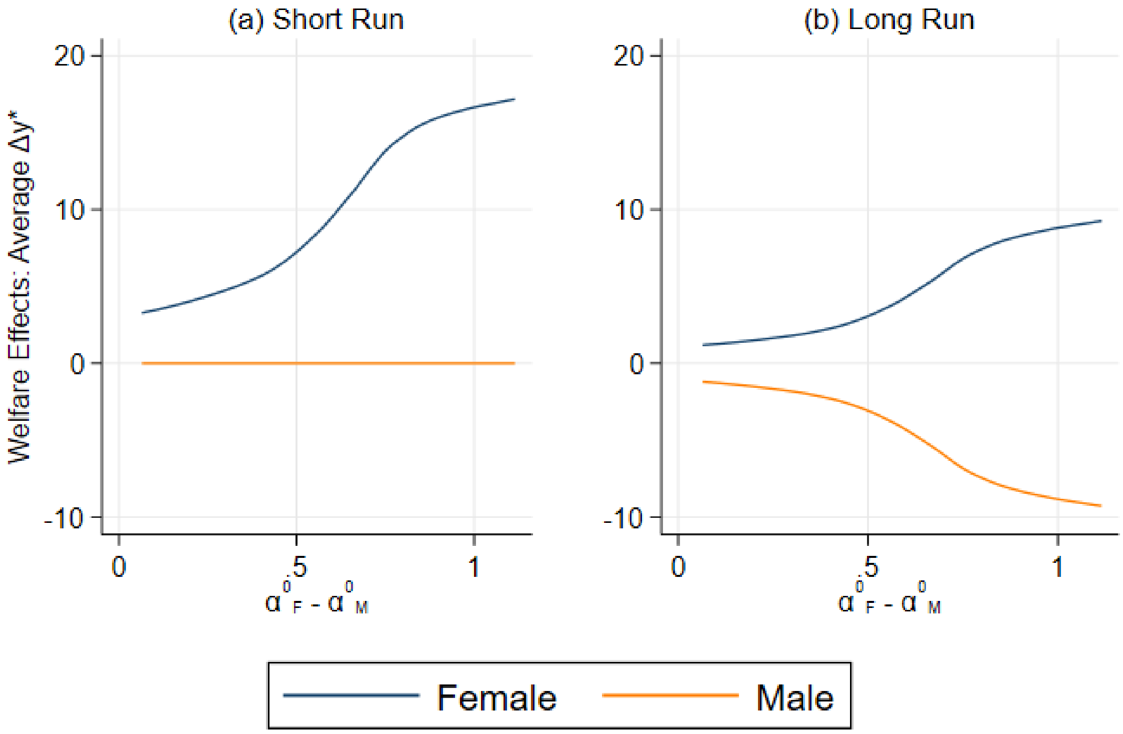

Notes. The figure shows the change in the value of chosen option y^* (which can be interpreted as a percentage change, because it is in units of log-wages) by the difference $\alpha_F^0 - \alpha_M^0$. For men, $y^* = \max\{y^0, y^1\}$. For women, $y^* = y_F^1 \mathbf{1}\{y_F^1 \geq y_F^0\} + y_M^0 \mathbf{1}\{y_F^1 < y_F^0\}$ (i.e. they make a decision based on y_F^0 but the true value of staying at home is y_M^0, since we interpret the difference $y_F^0 - y_M^0$ as the gender norm tax).